



**PILOT NON-CONFORMANCE TO
ALERTING SYSTEM COMMANDS DURING
CLOSELY SPACED PARALLEL APPROACHES**

Amy Ruth Pritchett and R. John Hansman

Aeronautical Systems Laboratory
Department of Aeronautics & Astronautics
Massachusetts Institute of Technology
Cambridge, MA 02139 USA

January 1997
ASL-97-2

Abstract

Cockpit alerting systems monitor potentially hazardous situations, both inside and outside the aircraft. When a hazard is projected to occur, the alerting system displays alerts and/or command decisions to the pilot. However, pilots have been observed to not conform to alerting system commands by delaying their response or by not following the automatic commands exactly. This non-conformance to the automatic alerting system can reduce its benefit. Therefore, a need exists to understand the causes and effects of pilot non-conformance in order to develop automatic alerting systems whose commands the pilots are more likely to follow.

These considerations were examined through flight simulator evaluations of the collision avoidance task during closely spaced parallel approaches. This task provided a useful case-study because the effects of non-conformance can be significant, given the time-critical nature of the task. A preliminary evaluation of alerting systems identified non-conformance in over 40% of the cases and a corresponding drop in collision avoidance performance. A follow-on experiment found subjects' alerting and maneuver selection criteria were consistent with different strategies than those used by automatic systems, indicating the pilot may potentially disagree with the alerting system if the pilot attempts to verify automatic alerts and commanded avoidance maneuvers. A final experiment found supporting automatic alerts with the explicit display of its underlying criteria resulted in more consistent subject reactions.

In light of these experimental results, a general discussion of pilot non-conformance is provided. Contributing factors in pilot non-conformance include a lack of confidence in the automatic system and mismatches between the alerting system's commands and the pilots' own decisions based on the information available to them. The effects of non-conformance on system performance are discussed. Possible methods of reconciling mismatches are given, and design considerations for alerting systems which alleviate the problem of non-conformance are provided.

This document is based on the thesis of Amy R. Pritchett submitted in partial fulfillment of the degree of Doctor of Science in Aeronautics and Astronautics at the Massachusetts Institute of Technology.

Table of Contents

LIST OF FIGURES	7
LIST OF TABLES.....	11
ACKNOWLEDGMENTS.....	13
1. INTRODUCTION & MOTIVATION	15
1.1 PILOT CONFORMANCE TO AUTOMATICALLY GENERATED COMMANDS	16
1.2 DEFINITION OF THE ALERTING TASK AND ROLES OF ALERTING SYSTEMS.....	17
1.3 CHARACTERISTICS OF COCKPIT ALERTING SYSTEMS	23
1.4 OVERVIEW OF THESIS	24
2. CLOSELY SPACED PARALLEL RUNWAY OPERATIONS.....	27
3. EXPERIMENT #1: SIMULATOR EVALUATION OF THE CHARACTERISTICS OF A COLLISION AVOIDANCE SYSTEM FOR PARALLEL APPROACHES.....	35
3.1 EXPERIMENT OBJECTIVES	35
3.2 EXPERIMENT DESIGN.....	36
3.3 EXPERIMENTAL RESULTS.....	46
3.4 SUMMARY OF RESULTS AND DISCUSSION.....	55
4. EXPERIMENT #2: SIMULATOR EVALUATION OF ALERTING AND AVOIDANCE STRATEGIES USED BY SUBJECTS	57
4.1 EXPERIMENT OBJECTIVES	57
4.2 EXPERIMENT DESIGN.....	58
4.3 EXPERIMENT RESULTS	63
4.4 SUMMARY OF RESULTS AND DISCUSSION.....	73
5. EXPERIMENT #3: DESIGN AND EVALUATION OF AN ALERTING SYSTEM TO ALLEViate PROBLEMS WITH PILOT NON-CONFORMANCE.....	75
5.1 EXPERIMENT OBJECTIVES	75
5.2 COMPARISON OF ALERT CRITERIA FOR CLOSELY SPACED PARALLEL APPROACHES.....	76
5.3 EXPERIMENT DESIGN.....	79
5.4 EXPERIMENT RESULTS	90
5.5 SUMMARY OF RESULTS AND DISCUSSION.....	97

6. PILOT CONFORMANCE TO AUTOMATICALLY GENERATED COMMANDS.....	99
6.1 PILOT INTERACTION WITH THE ALERTING SYSTEM, WITH AND WITHOUT CONFORMANCE	99
6.2 FACTORS CONTRIBUTING TO PILOT NON-CONFORMANCE.....	102
6.3 EFFECTS OF NON-CONFORMANCE ON ALERTING SYSTEM BENEFITS.....	109
6.4 THE IMPACT OF ALERTING SYSTEM DESIGN ON PILOT CONFORMANCE	111
7. CONCLUSIONS.....	115
7.1 COLLISION AVOIDANCE DURING CLOSELY SPACED PARALLEL APPROACHES.....	115
7.2 PILOT NON-CONFORMANCE TO ALERTING SYSTEM COMMANDS	116
7.3 RECOMMENDATIONS	117
REFERENCES	119
APPENDIX A CONTROL SIGNAL FLOW MODEL OF THE ALERTING SYSTEM.....	123
APPENDIX B MODELS OF THE HUMAN BEHAVIOR AT ALERTING TASKS	125
APPENDIX C DESCRIPTIONS OF ALERTING SYSTEM ROLES.....	139
APPENDIX D DESCRIPTION OF TRAFFIC ALERT AND COLLISION AVOIDANCE SYSTEM (TCAS II).....	149
APPENDIX E DESCRIPTION OF THE PRECISION RUNWAY MONITORING SYSTEM..	157
APPENDIX F MODELS OF POSSIBLE RECONCILIATION STRATEGIES.....	161

List of Figures

FIGURE 1.1 SCHEMATIC OF THE ALERTING AND AVOIDANCE TASKS	15
FIGURE 1.2 CONFORMANCE CRITERIA IMPLICIT IN SELECTION OF AN AVOIDANCE MANEUVER BY TCAS II	16
FIGURE 1.3 SUB-FUNCTIONS IN THE ALERTING & HAZARD RESOLUTION TASK	18
FIGURE 1.4 COMBINED HUMAN-AUTOMATION CONTROLLER	19
FIGURE 1.5 SCHEMATIC OF AN "EXECUTIVE, AUTOMATIC DECISION-MAKING" ROLE OF AN ALERTING SYSTEM... ..	21
FIGURE 1.6 SCHEMATIC OF A "SUPPORTING, AUTOMATIC HAZARD ASSESSMENT" ROLE OF AN ALERTING SYSTEM ..	22
FIGURE 2.1 ILLUSTRATION OF PARALLEL APPROACH PATHS	27
FIGURE 2.2 SCHEMATIC OF A BLUNDER BY ONE AIRCRAFT TOWARDS ANOTHER ON A PARALLEL APPROACH PATH ..	28
FIGURE 2.3 TIME TO COLLISION WITH DECREASING RUNWAY SEPARATION	28
FIGURE 2.4 DESCRIPTION OF THE 'KILL ZONE' CONCEPT	29
FIGURE 2.5 EXTENT OF THE KILL ZONE FOR DIFFERENT RUNWAY SEPARATIONS	30
FIGURE 2.6 POSITION OF THE KILL ZONE WITH DIFFERENT SPEEDS OF THE AIRCRAFT ON THE PARALLEL APPROACH ..	31
FIGURE 3.1 BASELINE TRAFFIC DISPLAY, SHOWING TRAFFIC ON PARALLEL APPROACH	37
FIGURE 3.2. ENHANCED EHSI TRAFFIC DISPLAY	39
FIGURE 3.3. TRAFFIC DISPLAY ON PFD	40
FIGURE 3.4 EXPERIMENT #1 TEST MATRIX	42
FIGURE 3.5 SIMULATOR CONTROL STATION AT TWO POINTS DURING A 'HAZARDOUS BLUNDER' SCENARIO	44
FIGURE 3.6 CRITERIA FOR JUDGING CONFORMANCE TO A TCAS SELECTED AVOIDANCE MANEUVER	46
FIGURE 3.7 FREQUENCY OF PILOTS' MANEUVERS NOT CONFORMING TO (OR NOT MATCHING) THE TCAS SUGGESTED MANEUVER, BY PROCEDURE	47
FIGURE 3.8 FREQUENCY OF NEAR-MISSES IN CONDITIONS WITH AND WITHOUT DISPLAY OF AVOIDANCE MANEUVERS, AND WITH AND WITHOUT PILOT CONFORMANCE TO THE MANEUVER	48
FIGURE 3.9 FREQUENCY OF EACH TYPE OF TURNING MANEUVER, COMPARED WITH THE DISPLAY OF, AND CONFORMANCE TO, TCAS AVOIDANCE MANEUVERS	49
FIGURE 3.10 FREQUENCY PILOT'S MANEUVER MATCHED TCAS MANEUVER, BY TRAFFIC DISPLAY	50
FIGURE 3.11. SUBJECTIVE AHP RATINGS OF THE DISPLAYS IN THE BASELINE EXPERIMENT FOR THE QUESTION "WHICH DISPLAY GIVES A CLEARER PICTURE OF THE TRAFFIC SITUATION?"	52
FIGURE 3.12. PERCENTAGE OF APPROACHES RESULTING IN AN EARLY AVOIDANCE MANEUVER, BY TRAFFIC DISPLAY	53
FIGURE 3.13 PILOT RANKINGS OF AUTOMATIC FUNCTIONS	54
FIGURE 4.2. EXPERIMENT #2 TEST MATRIX	62
FIGURE 4.3 HISTOGRAM OF LATERAL SEPARATION WHEN THE SUBJECTS REACTED	65
FIGURE 4.4 HISTOGRAM OF SECONDS REMAINING TO POINT OF CLOSEST APPROACH WHEN THE SUBJECTS REACTED	66
FIGURE 4.5 TIME REMAINING TO COLLISION WHEN THE SUBJECT REACTS, FOR DIFFERENT COLLISION GEOMETRIES	69

FIGURE 4.6. AIRCRAFT LATERAL SEPARATION AT THE SUBJECTS' REACTIONS, FOR DIFFERENT COLLISION GEOMETRIES	70
FIGURE 4.7 AIRCRAFT LATERAL SEPARATION AT SUBJECT REACTION TIMES WITH EACH DISPLAY	72
FIGURE 4.8 MEAN SQUARED ERROR IN THE SIDETASK FOR DIFFERENT SUBJECT WORKLOADS	73
FIGURE 5.1 PROBABILITY OF COLLISION CONTOURS.....	76
FIGURE 5.2 COMPARISON OF THE MIT AND NTZ ALERT CRITERIA.....	77
FIGURE 5.3 GEOMETRIC SIMPLIFICATION OF THE MIT ALERT CRITERIA	78
FIGURE 5.4 BASELINE TRAFFIC DISPLAY	81
FIGURE 5.5 NTZ ALERT CRITERIA SHOWN ON TRAFFIC DISPLAY	82
FIGURE 5.6 MIT ALERT CRITERIA SHOWN ON TRAFFIC DISPLAY.....	83
FIGURE 5.8 ILLUSTRATION OF THE FOUR DIFFERENT SCENARIOS	86
FIGURE 5.9. EXPERIMENT #3 TEST MATRIX	87
FIGURE 5.10 EXPANDED TEST MATRIX SHOWING NINE ALERT-DISPLAY COMBINATIONS	88
FIGURE 5.11 TIMING OF THE SUBJECTS REACTIONS IN THE LOW CONVERGENCE, HAZARDOUS BLUNDER SCENARIO, BY DISPLAY	91
FIGURE 5.12 TIMING OF THE SUBJECTS REACTIONS IN THE HIGH CONVERGENCE, HAZARDOUS BLUNDER SCENARIO, BY DISPLAY	91
FIGURE 5.13 PERFORMANCE OF SUBJECTS' REACTIONS WITH THE DIFFERENT DISPLAYS	92
FIGURE 5.14 NUMBER OF FREE RESPONSES IN EACH CATEGORY TO THE QUESTION "How Did The (AUTOMATIC) ALERTS CHANGE YOUR DECISIONS?"	94
FIGURE 5.15 DIFFERENCES IN TIME BETWEEN SUBJECTS' REACTIONS AND BOTH ALERTING CRITERIA, WHEN ALERTS ARE NOT GIVEN, BY DISPLAY	95
FIGURE 5.16 AVERAGE DIFFERENT IN TIME BETWEEN WHEN THE SUBJECT REACTED AND AN AUTOMATIC ALERT WAS TRIGGERED BY EACH TYPE OF ALERT CRITERIA, BY DISPLAY	96
FIGURE 5.17 PERCENTAGE OF 'LATE ALERTS' GIVEN IN THE HIGH CONVERGENCE RATE, HAZARDOUS BLUNDER	97
FIGURE 6.1 EXPECTED DEVELOPMENT OF CONTROL INPUTS TO THE AIRCRAFT WITH AN EXECUTIVE, DECISION-MAKING ALERTING SYSTEM AVAILABLE.....	100
FIGURE 6.2 ACTUAL DEVELOPMENT OF CONTROL INPUTS TO AIRCRAFT WHEN PILOT DOES NOT CONFORM	101
FIGURE 6.3 PILOT'S ACTIONS BASED UPON THEIR CONFIDENCE IN THE ALERTING SYSTEM AND THEIR OWN REASONING	106
FIGURE A.1 MODEL OF CONTROLLER FOR TIME-CRITICAL, DECISION-MAKING TASK	123
FIGURE A.2 THE INFORMATION PROCESSING SUB-FUNCTION	126
FIGURE A.3 THE HAZARD ASSESSMENT & ALERTING SUBFUNCTION.....	127
FIGURE A.4 THE DECISION MAKING AND SELECTION OF A HAZARD RESOLUTION SUB-FUNCTION.....	128
FIGURE A.5 SCHEMATIC OF THE CONSTRAINTS AND REQUIREMENTS ON EACH ALERTING TASK SUB-FUNCTION..	130
FIGURE B.1 HUMAN STRUCTURAL MODEL FOR COMPENSATORY CONTROL TASKS.....	133
FIGURE B.2 PATRICK'S MODEL OF THE DECISION PROCESS.....	137
FIGURE C.1 FULLY MANUAL CONTROL.....	140

FIGURE C.2 AUTOMATIC INFORMATION PROCESSING.....	141
FIGURE C.3 AUTOMATIC HAZARD ASSESSMENT AND ALERTING.....	142
FIGURE C.4 AUTOMATIC DECISION MAKING.....	144
FIGURE C.5 FULLY AUTOMATIC CONTROL.....	145
FIGURE D.1 SCHEMATIC OF TCAS II COMPONENTS.....	149
FIGURE D.2 TCAS II SENSOR INFORMATION.....	150
FIGURE D.3 TCAS ALERTING CRITERIA, BASED ON INTER-AIRCRAFT RANGE AND RANGE-RATE.....	151
FIGURE D.4 SCHEMATIC OF MANEUVERS EVALUATED (CLIMB SENSE).....	152
FIGURE D.5 TCAS II COCKPIT TRAFFIC DISPLAY, INTEGRATED ONTO NAVIGATION DISPLAY	153
FIGURE D.6 TCAS PRESENTATION OF AVOIDANCE MANEUVER ON PRIMARY FLIGHT DISPLAY.....	154
FIGURE E.1 NON-TRANSGRESSION ZONE CRITERIA USED BY PRM	158
FIGURE E.2 SCHEMATIC OF FALSE ALARMS CAUSED BY AIRCRAFT DEVIATIONS INTO OR TOWARDS THE NTZ DURING NOMINAL APPROACHES.....	159
FIGURE E.3 TIME LINE OF PRM COLLISION AVOIDANCE SEQUENCE.....	160
FIGURE F.1 PARALLEL REASONING OF THE PILOT AND ALERTING SYSTEM FOR THE ALERTING SUB-TASK; FINAL ALERTING DECISION BASED ON RECONCILIATION OF EACH.....	161
FIGURE F.2 SERIAL METHOD OF RECONCILIATION.....	162
FIGURE F.3 RECONSTRUCTION OF ALERTING SYSTEM'S LOGIC.....	163

List of Tables

TABLE 2.1 CATEGORIES OF ALERTS	33
TABLE 2.2 TRADE-OFF BETWEEN ALERTING ACCURACY AND AVOIDANCE MANEUVER EFFECTIVENESS.....	33
TABLE 4.1 CORRECTNESS OF SUBJECTS' REACTIONS.....	64
TABLE 4.2 FREQUENCY OF AVOIDANCE MANEUVERS SELECTION.....	66
TABLE 4.3 POTENTIAL TIMES TO COLLISION WITH DIFFERENT SCENARIO VARIABLES.....	67
TABLE 4.4 CORRECTNESS OF THE SUBJECTS' REACTIONS TO THE HIGH AND LOW CONVERGENCE INTRUSIONS.....	68
TABLE 4.5 CORRECTNESS OF SUBJECT REACTIONS WITH THE DIFFERENT DISPLAYS.....	71
TABLE 6.1 COMPARING PILOT'S LOW CONFIDENCE IN ALERTING SYSTEMS WITH THE 7 CAUSES OF TRUST IDENTIFIED BY SHERIDAN.....	107
TABLE B.1 EXAMPLES OF TASK BREAKDOWNS USED IN OTHER STUDIES.....	135

Acknowledgments

This work was supported by NASA Ames Research Center under Grant NAG 2-716. The authors wish to thank their technical monitor, Dr. Kevin Corker, for his contributions to the project.

The authors also wish to thank the numerous pilots and subjects who donated their time to participate in the flight simulator experiments.

1. Introduction & Motivation

Cockpit alerting systems provide an automatic means to assess hazards, evaluate if an alert is required to cue an action and, with some alerting systems, decide upon methods to prevent or resolve the hazardous condition, such as the change in trajectory shown schematically in Figure 1.1. The hazards evaluated by the alerting system may be internal to the aircraft (e.g. an engine failure) or they may include external conditions (e.g. a collision with another aircraft or the ground). The assessment of the hazard may be currently affecting the aircraft (e.g. again, an engine failure), or it may be projected to occur in the near future (e.g. again, an aircraft collision avoidance system).

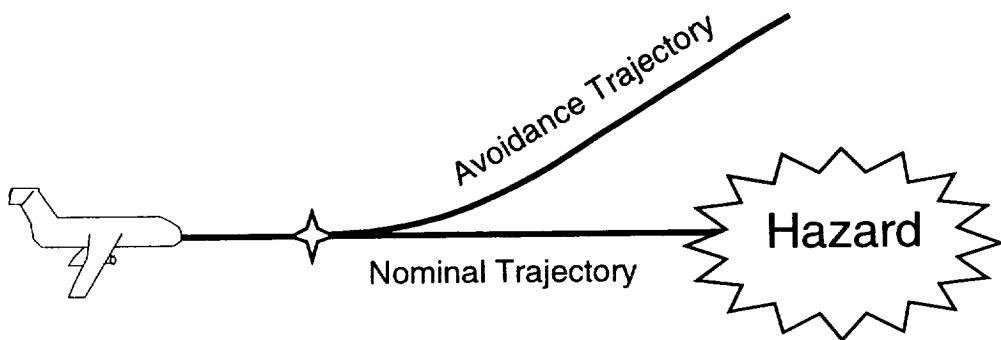


Figure 1.1 Schematic of the Alerting and Avoidance Tasks

Alerting systems with increasing complexity and sophistication are being developed. These alerting systems can be given executive roles (i.e. designed with the implicit assumption that their commands will be executed quickly and precisely by the pilot.) For example, with the Traffic alert and Collision Avoidance System (TCAS), conflicting aircraft agree on their respective directions for an avoidance maneuver. The avoidance maneuvers are then calculated with the assumption that one aircraft will climb, the other descend. Either pilot's decision to not conform to the alerting system represents an unexpected variation with possibly positive, possibly negative effects. (RTCA, 1983)

This thesis examines the causes and implications of pilot non-conformance to the automatic commands displayed by an alerting system.

1.1 Pilot Conformance to Automatically Generated Commands

Implicit in the design of an executive alerting system is the assumption that pilots will conform to the alerting system commands. For example, as shown in Figure 1.2, the current Traffic alert and Collision Avoidance System (TCAS II) assumes, once an alert has been issued, that a pilot will have a response delay of no more than five seconds, followed by a quarter-G pull-up (or push-over, if the commanded maneuver is a descent), and that the commanded vertical speed will be followed for as long as it is displayed. (RTCA, 1983)

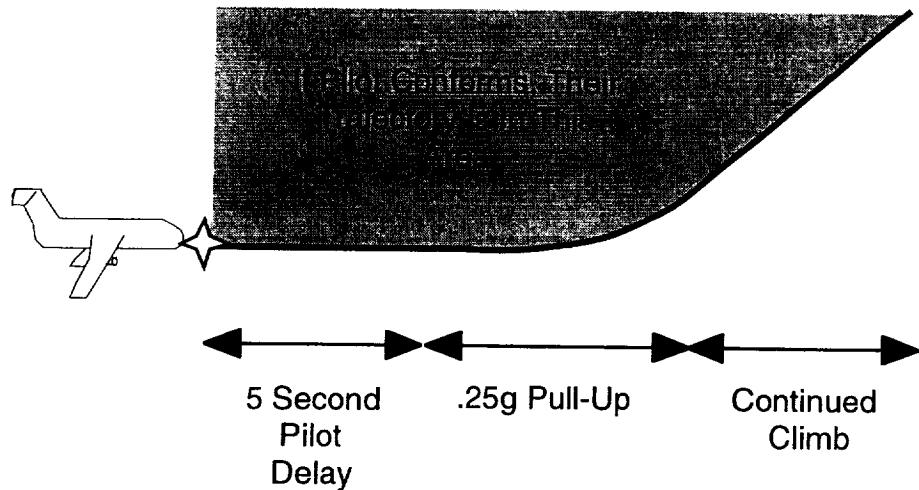


Figure 1.2 Conformance Criteria Implicit in Selection of an Avoidance Maneuver by TCAS II

The assumption that pilots will conform is not always valid. Studies of currently operational alerting systems have identified non-conformance situations where pilots have delayed in responding to automatic alerts, or have executed different resolutions to the hazard than commanded by the automatic system.

This non-conformance may be intentional or unintentional. In cases of unintentional non-conformance, pilots are attempting to follow the alerting system's commands, but can not because of conflicting concerns or because of misinterpretation of the alerting system's commands. Concerns about unintentional non-conformance are typically addressed during the design and testing of an alerting system. (e.g. Mårtensson, 1995)

This thesis focuses instead on intentional non-conformance to the alerting system's commands. Cases of intentional non-conformance have been found with executive alerting systems. For

example, pilot questionnaires on the use of TCAS II reported pilots intentionally did not follow commanded avoidance maneuvers in 24.7% of the cases where alerts and commands were given. (Ciemier et al, 1993) A similar tendency was noted during a survey in which pilots reported delaying their response to an alert from the Ground Proximity Warning System (GPWS), and executing a less severe avoidance maneuver than the proceduralized response that is expected. (DeCelles, 1991)

The design and evaluation of executive alerting systems generally assumes pilots will have a variable reaction time and will execute the alerting system commands with some level of precision. However, instances of intentional non-conformance appear to indicate a separate, unanticipated variation, where pilots are additionally performing the cognitive task of assessing the situation and the automatic command's validity, and then possibly electing to use a different resolution to the hazard than given by the alerting system.

Non-conformance, therefore, can dramatically change the actions taken by pilots and the resulting performance attained. When the pilots have more information about the current situation or have a better understanding of the task, non-conformance can result in higher performance than expected from the alerting system. For example, pilot non-conformance to TCAS II alerts can be beneficial in instances where they have visual contact with the other aircraft and, through radio communications, know the other aircraft is continuing on a safe trajectory.

However, many alerting systems are intended for conditions where pilots may not have enough information, enough time, or enough free attention to use high performance decision strategies. In these cases, pilot non-conformance may reduce the positive benefit in performance expected from the availability of the automatic commands, and does not allow the alerting system to relieve pilots of the work load of the decision-making tasks.

1.2 Definition of the Alerting Task and Roles of Alerting Systems

The simplest alerting systems serve as a trigger for pilots to perform their own assessment of all the steps in the alerting task. For example, a stall warning indicates to pilots a proximity to a stalled attitude. The pilots may evaluate the situation and assess if a hazard exists, if action is required, and, if so, what that action should be.

More sophisticated alerting systems, however, can take on a variety of roles. The many possible roles of an alerting system can make comparisons between systems difficult; no commonly accepted terminology is available to succinctly describe the role of an alerting system. This section provides definitions of the various roles of alerting systems based on two criteria.

The first criterion compares the elements of the alerting task performed by the pilot to those performed by the alerting system. An alerting task can be broken down into the four serial sub-functions shown in Figure 1.3:

- *Information Processing* This sub-function transforms the inputs from the sensors to generate the relevant set of current values required for the hazard assessment. These operations may include fusing several different sensor values, and filtering and estimating noisy or incomplete sensor data.
- *Hazard Assessment & Alerting* This sub-function transforms knowledge of the current situation into an assessment of current or future hazards, and decides if the hazards require an alert to cue an off-nominal action.
- *Decision Making & Selection of Resolution Method* Once an alert has been issued, this sub-function selects a resolution to the hazard.
- *Control Actuation* This sub-function manipulates the aircraft controls such that the resolution method is carried out.

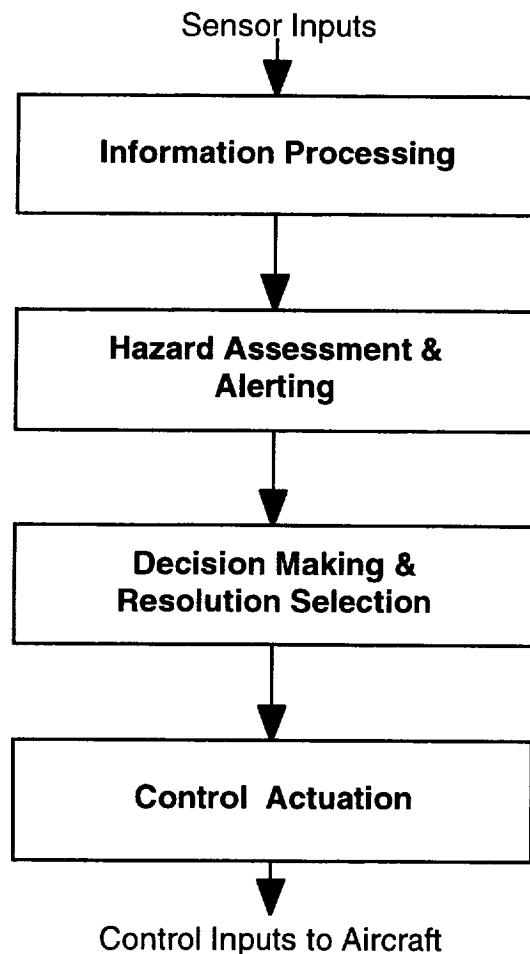


Figure 1.3 Sub-Functions in the Alerting & Hazard Resolution Task

When both the pilot and the alerting system can perform any of these sub-functions, several different sequences can be used to generate the commands to the aircraft, as shown in Figure 1.4. Five distinct roles of alerting systems can be defined using this criterion:

- *Fully Manual Control* The pilot is responsible for all sub-functions in the alerting task (i.e. no alerting system is available).
- *Automatic Information Processing* The alerting system provides a display of current state information which it has derived through filtering, combination and estimation based on the current sensor values (e.g. a collision avoidance system can provide detailed, processed information to the pilot about the position of other aircraft, leaving the pilot with the alerting and avoidance maneuver selection tasks).

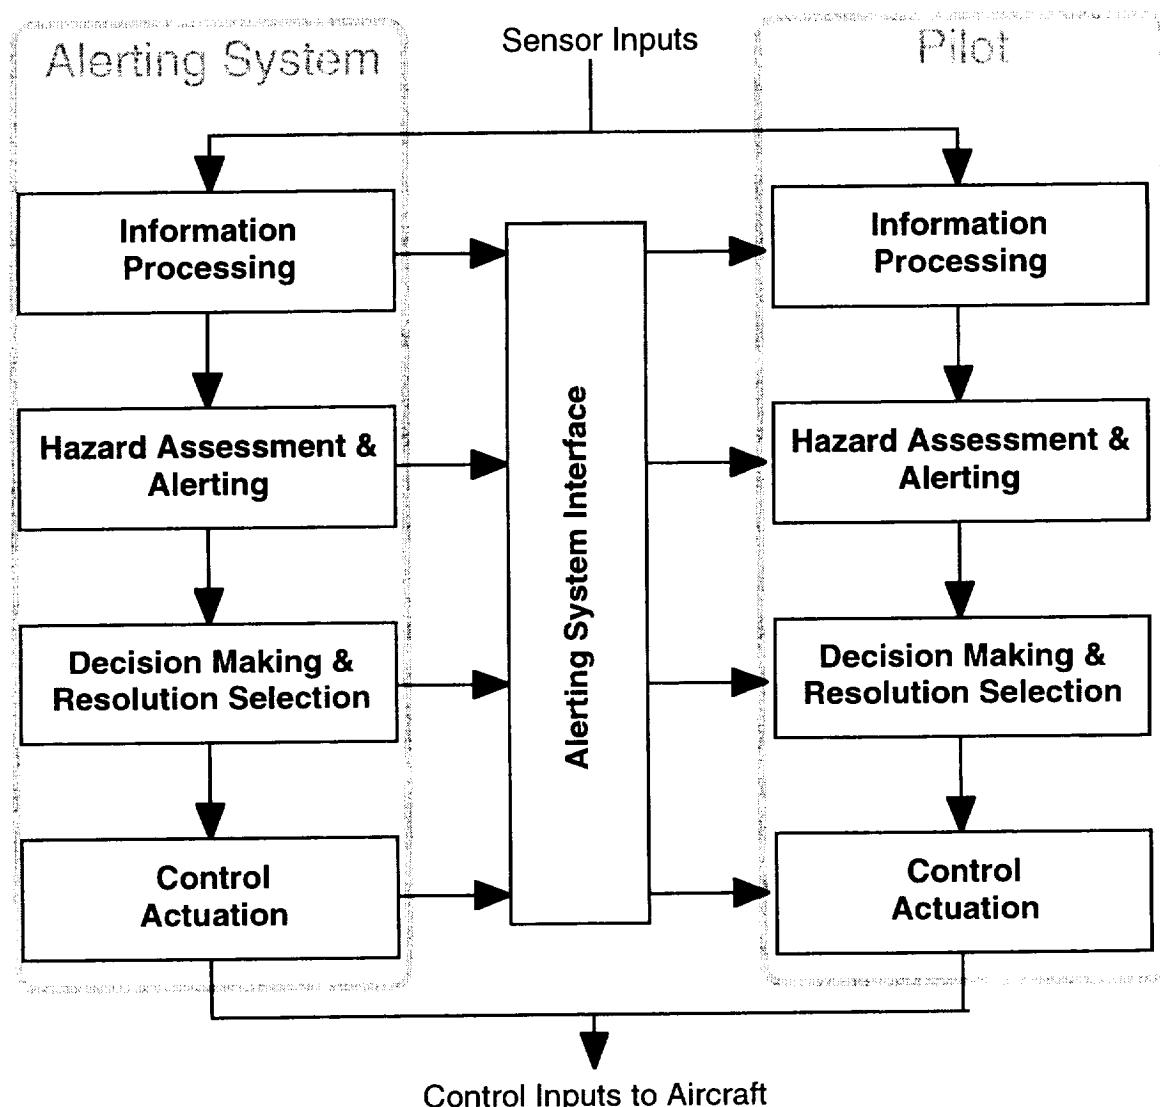


Figure 1.4 Combined Human-Automation Controller

- *Automatic Hazard Assessment & Alerting* The alerting system provides the pilot with a display of the projected hazard levels and/or discrete warnings and alerts. The alerting system may also provide current state information to the pilots, or it may be assumed the pilots' need for current state information is relieved by the presentation of the hazard information (e.g. a collision avoidance system may generate aural alerts indicating a potential collision hazard; a display of the current traffic display may be included).
- *Automatic Decision Making* The alerting system provides the pilot with possible resolutions to the hazard and/or command decisions. The alerting system may also provide current processed information and alert information to the pilot (e.g. the TCAS II system provides all three levels of information: traffic displays, alerts, and commanded avoidance maneuvers for the pilot to follow).
- *Fully Automatic* The alerting system, when a resolution is required, takes control and executes the resolution to the hazard. The alerting system may also provide current processed information, alert information, and a display of the resolution being executed to the pilot (e.g. an automatic system in the SR-71, upon detecting an engine failure, shuts down the other engine to prevent development of an uncontrollable yaw from asymmetric thrust).

The second criterion provides finer definition of the authority of the alerting system in each sub-function of the alerting task. If the alerting system is designed to replace pilots at a sub-function, it can be categorized as 'Executive'. The alerting system can instead be categorized as 'Supportive' if it is designed to support pilots at performing the sub-function through presentation of additional information derived through constant vigilance, additional computation, memory of recent data values, and databases which the pilot may otherwise not be able to perceive or calculate. Together, these two criteria provide a succinct, relatively descriptive method of describing the typical roles of alerting systems.

An 'Executive' alerting system is often used when the sub-function is thought to be too difficult or too time-critical for the pilot to be able to perform, or when the sub-function is thought too tedious and mundane to require the pilot's attention. The TCAS II system described previously is an example of an alerting system with an executive role.

The development of the aircraft commands in an 'Executive, Automatic Decision-Making' alerting-system is shown in Figure 1.5. The expected control signal evolves through the alerting system until it is displayed for the pilot to execute. While information from the previous sub-functions may be shown to the pilot, it is not expected to change the actions of the pilot. In addition, the display of information from one sub-function is assumed to eliminate the need for the pilot to perform that sub-function.

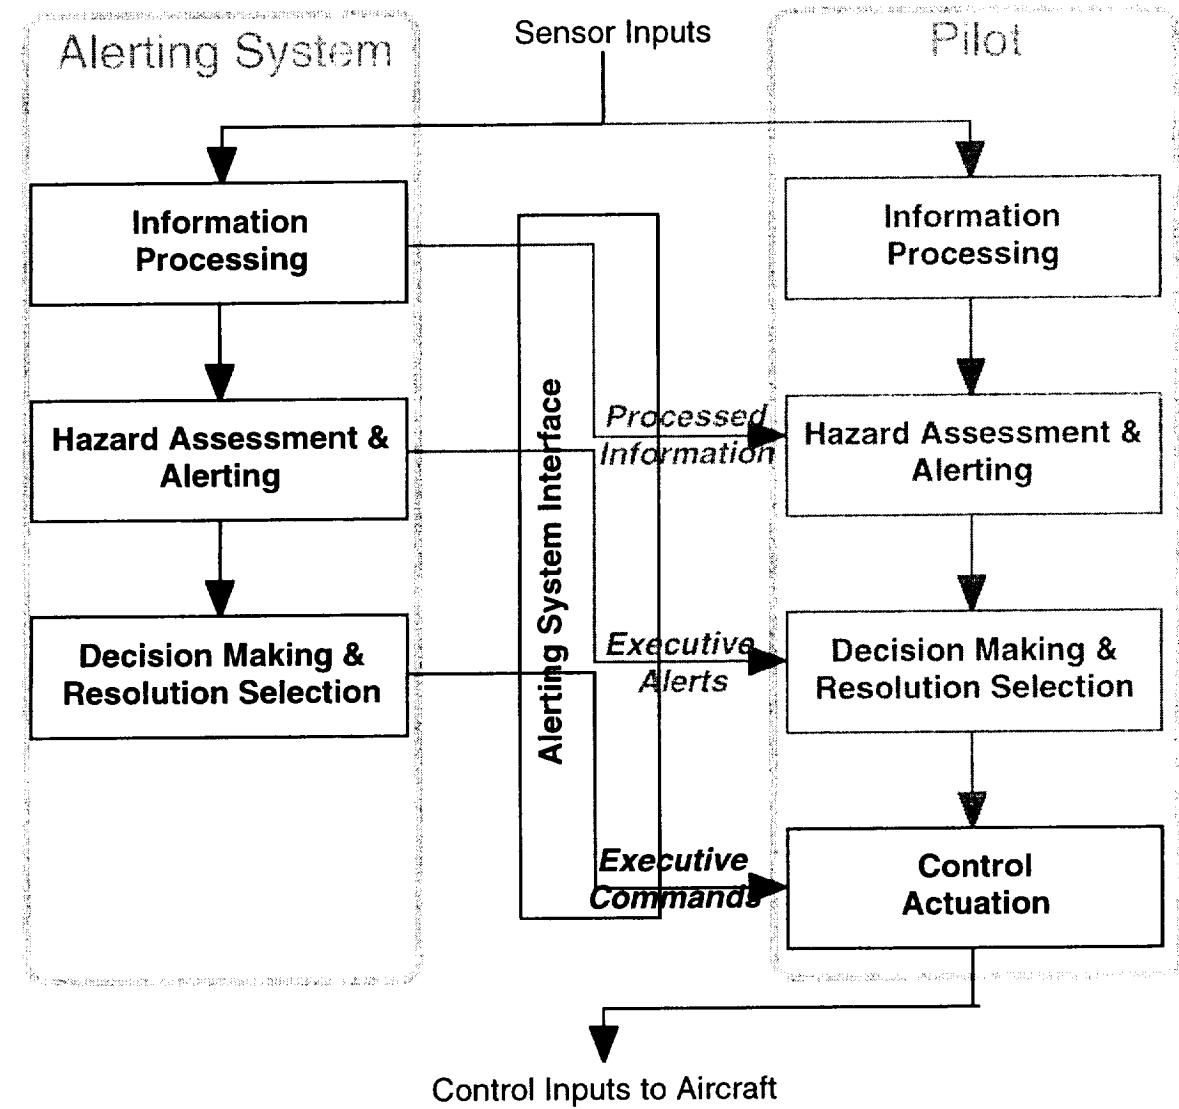


Figure 1.5 Schematic of an ‘Executive, Automatic Decision-Making’ Role of an Alerting System

A ‘Supportive’ type of alerting system uses the computational power and storage capacity of the automatic system to provide pilots with information they might not otherwise be able to incorporate into their decision-making. The indication given by Flight Management Systems of projected fuel remaining at the destination, with an alerting message if the projected fuel is too low, is an example of a supportive alerting system.

For example, Figure 1.6 shows an alerting system in a ‘Supporting Automatic Hazard Assessment’ role. In this case, the alerting system displays the current state information and predicted hazard levels, leaving the pilot to decide upon the need for an alert and to select a resolution. The information displayed by the alerting system for each sub-function is modeled as feeding into the pilot’s execution of the same sub-function, instead of replacing the pilot’s outputs, illustrating its supporting role.

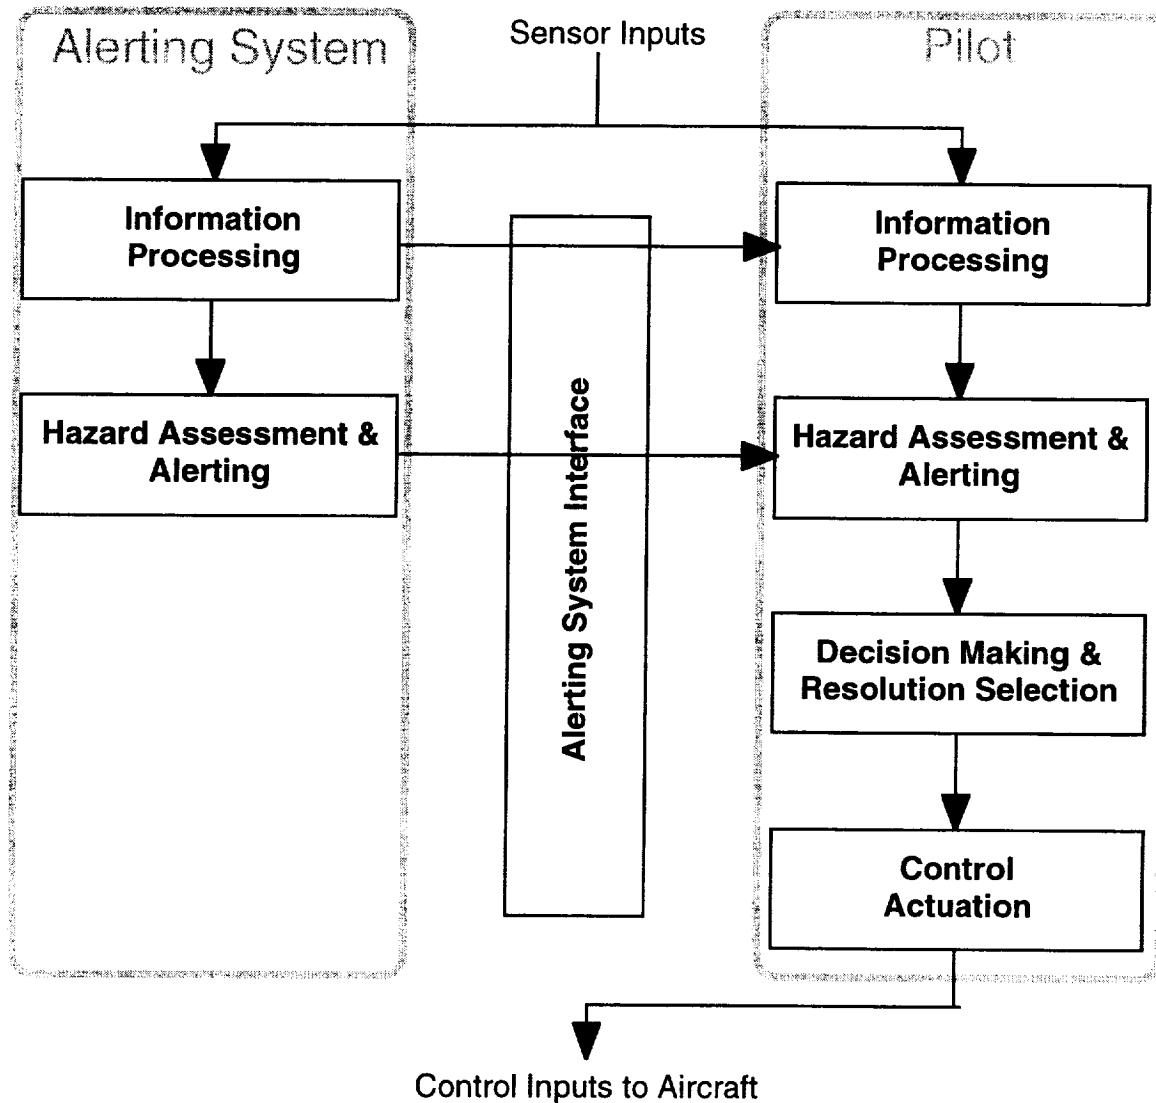


Figure 1.6 Schematic of a “Supporting, Automatic Hazard Assessment” Role of an Alerting System

This section has briefly defined the sub-functions in the alerting task; Appendix A contains a more detailed description of their internal processes.

The model in this section has divided the alerting task into four serial sub-functions; Appendix B compares the break down of the alerting task given in this section to models and concepts used in behavioral studies of human operators at similar tasks.

The definitions in this section have provided definitions for the most common roles of an alerting system. Other possible roles of automatic systems, and finer distinctions of automatic roles, are given in Appendix C.

1.3 Characteristics of Cockpit Alerting Systems

The simplest cockpit alerting systems, such as a stall warning, provide indications to pilots of a well-defined condition. These alerting systems tend to use a single sensor input, and provide a single, unambiguous indication. The performance of this simple type of alerting system is measured simply as whether it recognizes an anomalous sensor reading.

More sophisticated alerting systems are being developed to meet a variety of demands. These demands include: the requirement to monitor more, and more complex, cockpit systems; the trend towards more ‘optimal’ or economical air traffic management; the desire for higher safety margins; and the requirement to operate within tighter constraints, such as closer proximity to other aircraft, without a reduction in safety.

The performance of these more sophisticated alerting systems may be measured by the system attributes as given by the aircraft states and external conditions. Performance specifications for these systems may include a variety of minimal criteria (e.g. specifying an aircraft miss distance of at least 500 feet) and optimizing criteria (e.g. specifying an avoidance maneuver generate the least deviation from the intended flight path).

To achieve these types of performance specifications, cockpit alerting systems are performing increasingly sophisticated tasks. These tasks can have any of the following traits:

- Processing of several sensor inputs (e.g. in aircraft collision avoidance tasks, the trajectory of both the own aircraft and other aircraft must be projected based on data from a variety of sources);
- Dependence on multiple agents following a specified protocol (e.g. in a traffic conflict between two aircraft equipped with TCAS II, the two systems negotiate divergent directions for each aircraft’s avoidance maneuver, each of which needs to be followed for the full benefit of the system to be realized);

- Incomplete knowledge about the nominal behavior of the relevant agents (e.g. in aircraft collision avoidance tasks, the alerting system may not have complete knowledge about the nominal trajectory of both the own aircraft and the other aircraft);
- Incomplete knowledge about the efficacy of resolutions to the hazard (e.g. in aircraft collision avoidance tasks, the performance of an avoidance maneuver is dependent upon the maneuver the other aircraft will follow, which may not be known); and
- Multi-objective performance specifications, for which the relative importance of the different performance specifications may be arbitrary and whose elements may be contradictory (e.g. in aircraft collision avoidance tasks, a trade-off must be made between the severity of an avoidance maneuver and the resulting aircraft separation).

Alerting systems which have these traits also have a commensurate increase in the complexity of their underlying logic. Because this logic may be considered too sophisticated for pilots to perform unassisted, these alerting systems may be given executive roles. However, the resulting alerting system commands may appear non-intuitive to the pilot and non-conformance may be an issue.

1.4 Overview of Thesis

This thesis analyzes factors contributing to pilot non-conformance to automatic alerting systems. As a case study, alerting systems for collision avoidance during closely spaced parallel approaches are examined. As outlined in Chapter 2, this specific task is highly time-critical, and the alerting and decision-making criteria used by prototype traffic alerting and collision avoidance systems are sometimes non-intuitive, creating the opportunity for mismatches to occur between decisions made by the automatic alerting system and by pilots.

Chapter 3 describes a simulator experiment which examined the use of cockpit alerting systems and cockpit traffic displays by pilots during parallel approaches. Better collision avoidance performance and higher pilot opinions identified a need for an automatic system which presents alerts and avoidance maneuvers. However, pilots were found to not conform to the automatically calculated avoidance maneuvers, as they were displayed, in 40% of the approaches, with a resulting lower performance at avoiding the intruding aircraft.

In Chapter 4, a simulator experiment is described which tested the types of alerting and avoidance maneuver strategies favored by subjects during parallel approaches when an alerting system is not available. Subjects' decisions were consistent with simple alerting and decision making strategies, which are susceptible to certain types of collision trajectories. These simple

strategies can conflict, under certain conditions, with the more sophisticated strategies suitable for an alerting system, creating the potential for mismatches.

Chapter 5 details an experiment testing the effect of explicitly displaying the alert criteria used by sophisticated strategies during parallel approaches. Subjects were found to have reactions at times more consistent with sophisticated decision strategies, and to believe automatic systems using these strategies, when the strategies' decision criteria are presented in manner allowing the subjects to evaluate them easily.

Chapter 6 analyzes the results of the experiments and discusses the general causes and effects of pilot non-conformance. Pilots' confirmation of the alerting system's commands is hypothesized to represent a lack of confidence in the alerting system. This confirmation of the alerting system's commands can produce an unexpected time delay in the pilots' responses; when pilots also disagree with the alerting system they may execute a different resolution to the hazard. These actions indicate the presence of an executive alerting system does not necessarily reduce pilot workload, nor does it guarantee higher performance resolutions to hazards will be recognized by pilots. These issues raise considerations in the design of alerting systems.

Chapter 7 provides a summary and outlines the major conclusions of the thesis.

Appendices A, B and C elaborate on the definitions and descriptions of alerting tasks and the roles of alerting systems given in this example.

Appendices D and E provide background information about the Traffic alerting and Collision Avoidance System (TCAS) and Precision Runway Monitor (PRM).

Appendix F details possible methods a pilot may use to reconcile their decisions with alerting system commands.

2. Closely Spaced Parallel Runway Operations

This chapter briefly describes the task of collision avoidance during closely spaced parallel approaches, which will be used throughout the remainder of this thesis as a case study. Several major airports around the United States have, or plan to have, closely-spaced parallel runways. When aircraft are landing on both runways simultaneously the aircraft must fly close to each other during their parallel landing approaches, as shown in Figure 2.1. During Visual Meteorological Conditions (VMC), the responsibility for collision avoidance is given to the pilots, who are to maintain visual contact with each other. However, during Instrument Meteorological Conditions (IMC), the responsibility is currently given to air traffic controllers. Current technology limits independent parallel approaches to runways spaced 4300 feet or more apart, 3000' feet with the Precision Runway Monitor (PRM), which uses specialized ground based radar and a dedicated controller. The use of new technologies to reduce this minimum separation would allow airports to effectively maintain their high VMC landing capacity in IMC.

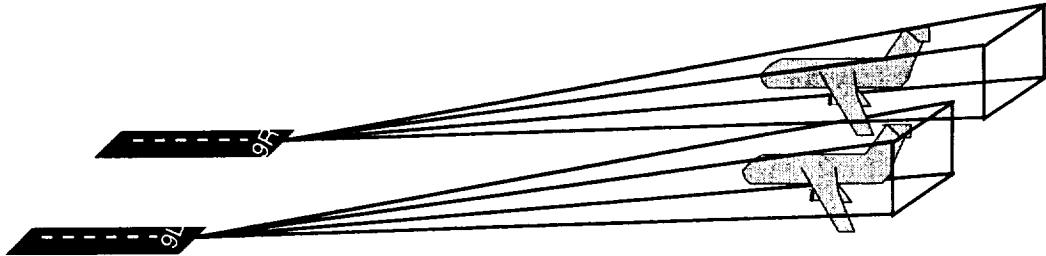


Figure 2.1 Illustration of Parallel Approach Paths

Several issues need to be resolved to enable closely spaced parallel approaches. Aircraft normally waver about their approach path; the amplitude of this wavering has been found to increase with distance from the runway threshold, and can be large enough to limit runway spacing and/or the distance from the threshold where the approach is intercepted by the aircraft. During a nominal approach, the wake vortex of an aircraft may blow into a parallel approach path before dissipating; wake vortex from an aircraft straying from its approach path may also affect aircraft on a parallel approach, and may limit the types of missed approach procedures used. Finally, for maximum runway throughput to be realized, a mechanism must be in place to space the aircraft onto the parallel approaches with the greatest efficiency. (Owen, 1993; Corjon & Poinsot, 1996; Koczo, 1996)

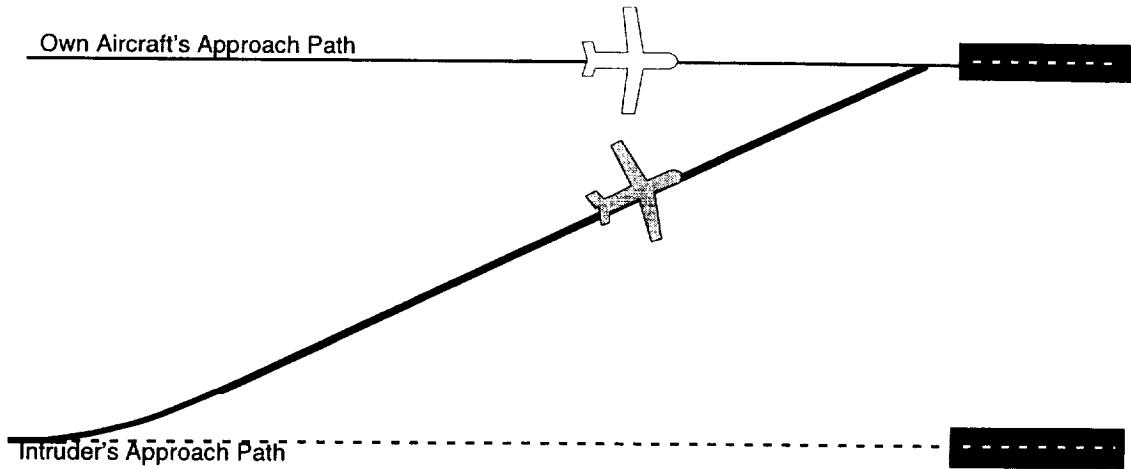


Figure 2.2 Schematic of a Blunder by One Aircraft Towards Another on a Parallel Approach Path (Top-Down View)

This thesis focuses on the task of collision avoidance during closely spaced parallel approaches. A schematic of a possible blunder by an intruder aircraft towards another aircraft is shown in Figure 2.2. Ensuring adequate aircraft separation in such a situation is a difficult task. The aircraft are closer together than during any other airborne phase of flight. As shown in Figure 2.3, a collision could potentially occur within seconds. Decisions to alert and select an avoidance

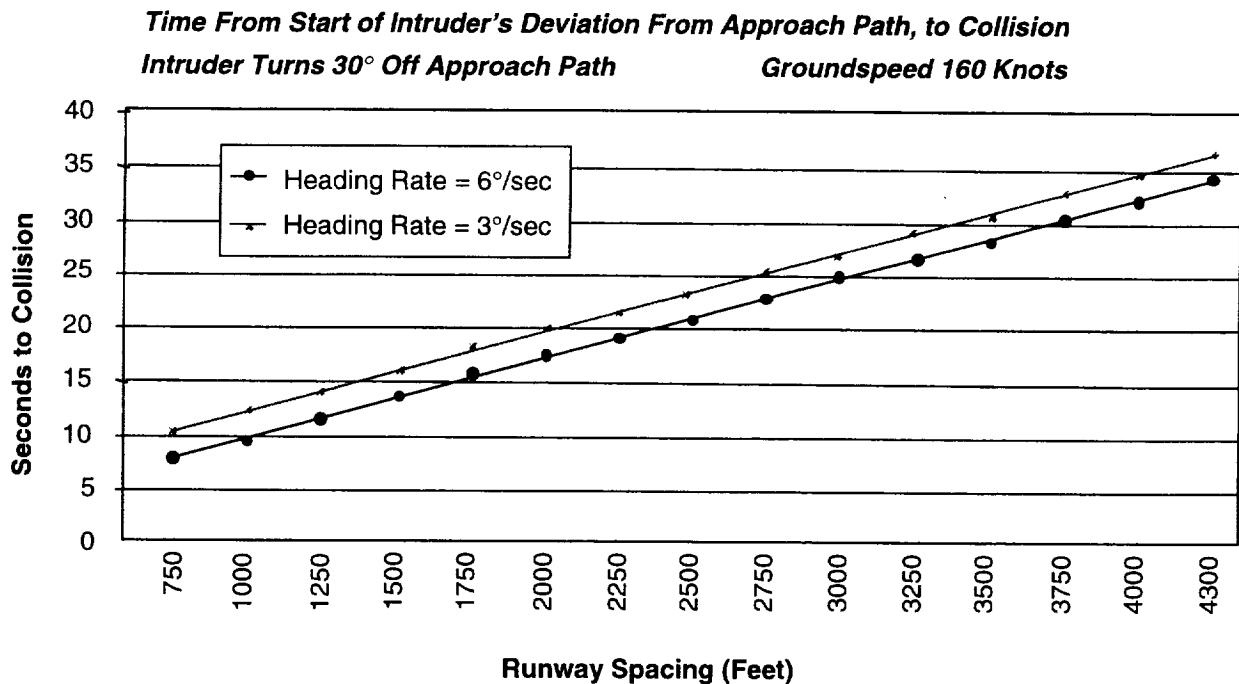


Figure 2.3 Time to Collision With Decreasing Runway Separation

maneuver can not assume the intruding aircraft will follow a predictable trajectory, and can not wait for more stable information about the intruder's future trajectory. Despite these difficulties, the collision avoidance system's decisions must provide adequate separation for any reasonable intruder aircraft trajectory while minimizing the number of nuisance alerts, which force unnecessary missed approaches and thus reduce airport capacity.

The constrained geometry created by parallel approaches results in a limited set of relative positions between the aircraft from which a collision hazard could be created. For any 'own aircraft', this creates a 'Kill Zone' -- a range of possibly dangerous relative positions -- for aircraft on the parallel approach, as shown in Figure 2.4. If an aircraft on a parallel approach is in front or behind the Kill Zone, any reasonable maneuver will steer it in front or behind of the own aircraft by a sufficient safety margin.

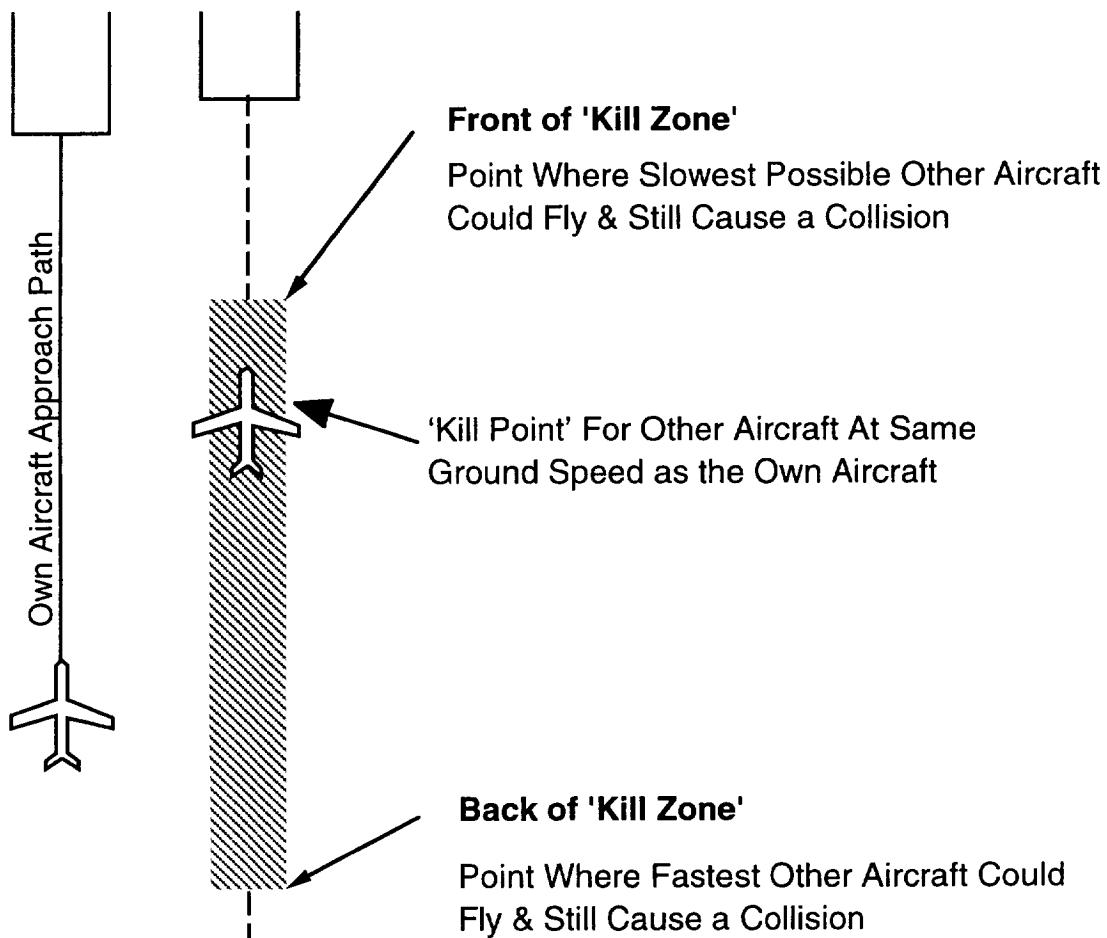


Figure 2.4 Description of the 'Kill Zone' Concept

The extent of the Kill Zone is affected by several factors. As shown in Figure 2.5, the Kill Zone becomes smaller with smaller runway separation. Also, as shown in Figure 2.6, the position of the Kill Zone is strongly dependent on the other aircraft's relative speed. These figures assumed specific limits on intruder heading and bank; if more severe maneuvers by the intruder are possible, the Kill Zone increases in size.

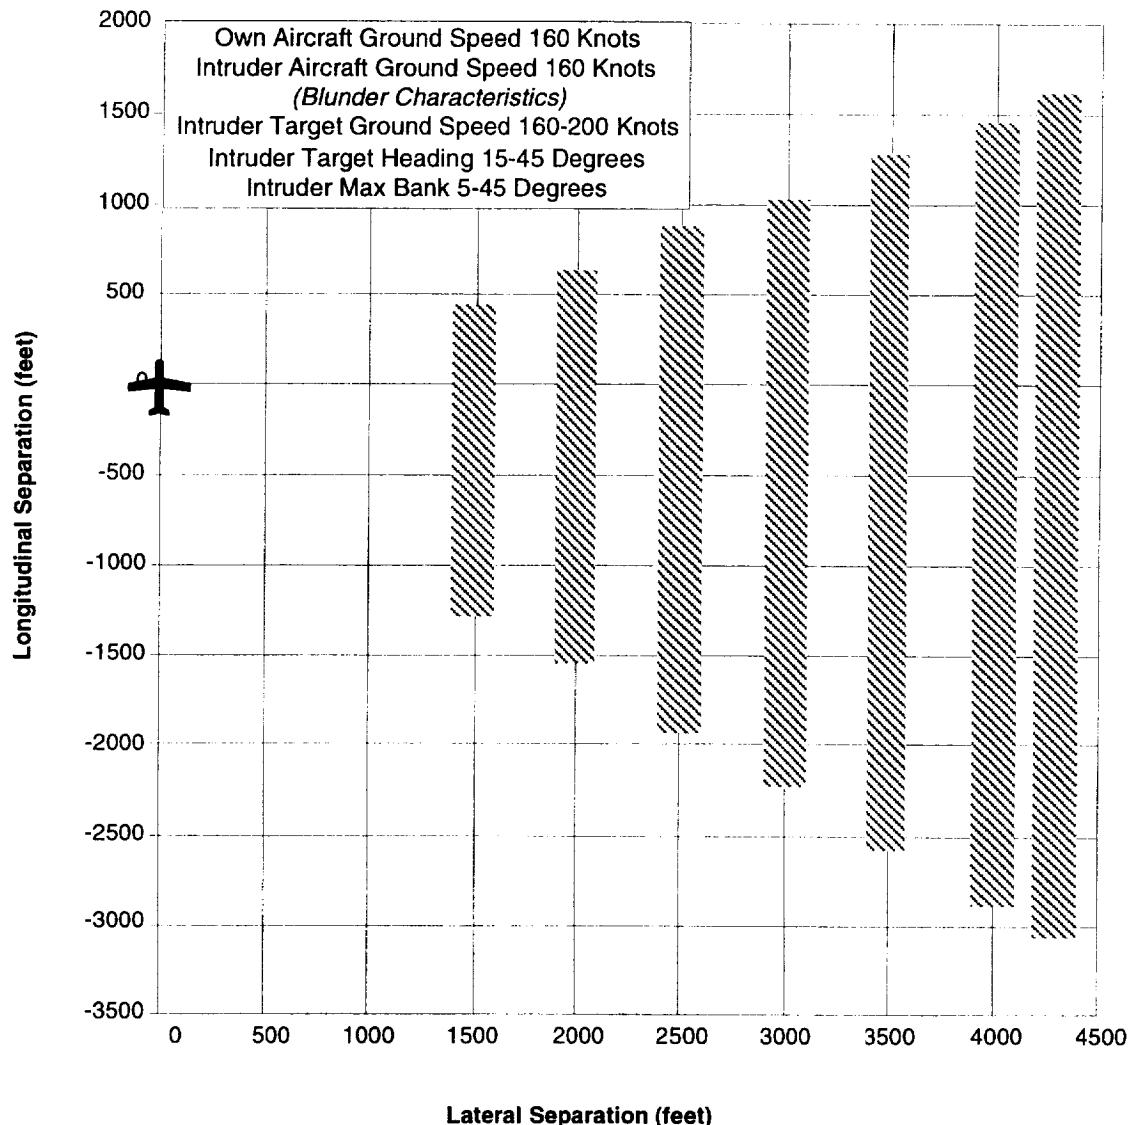


Figure 2.5 Extent of the Kill Zone for Different Runway Separations

The existence of the Kill Zone has several implications. First, the relative position of the Kill Zone may be non-intuitive. For example, two parallel aircraft with the same ground speed and directly abeam each other, especially at lower runway separations, are generally considered to represent a hazardous situation; in fact, extreme maneuvers would be required for a collision to occur from this relative position when the aircraft have the same ground speed. Second, the Kill Zone highlights the difficulties in deciding upon the need for a collision alert during parallel approaches, as the Kill Zone during parallel approaches is much smaller and well-defined than during any other phase of flight, representing the unusual nature of this type of operations.

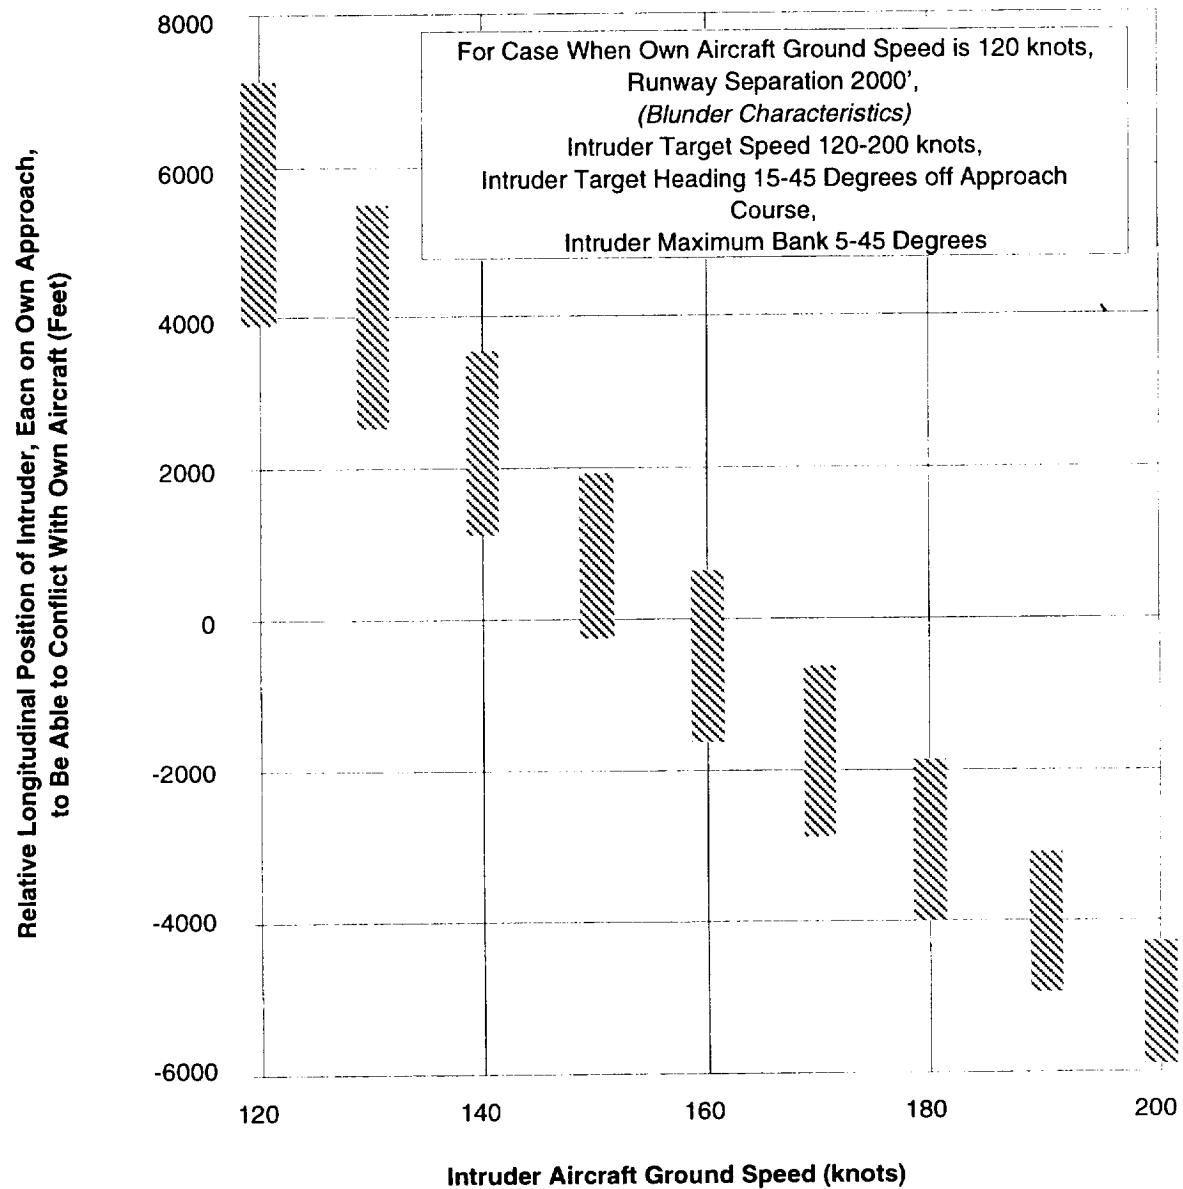


Figure 2.6 Position of the Kill Zone With Different Speeds of the Aircraft on the Parallel Approach

Current collision avoidance systems are not capable of ensuring adequate separation for closer runway spacings. Both the ground-based PRM system and the airborne Traffic alert and Collision Avoidance System (TCAS II) have been shown to have limitations preventing their use for close runway spacings. Appendices E and F describe the relative merits of PRM and TCAS II. Several studies are examining the use of new technologies, such as the ability to cross-link aircraft state data and specialized alerting logic, in order to enable a cockpit collision avoidance system specifically for this flight condition. (e.g. Carpenter & Kuchar, 1996; Koczo, 1996)

The primary goal of a collision avoidance system for this phase of flight is the requirement that the aircraft always maintain an acceptable separation. This thesis defines an acceptable separation as a miss distance of at least 500 feet slant range separation between the aircraft centers of mass, the same criteria used in other studies. (Carpenter & Kuchar, 1996; Wong, 1993) Other goals also require consideration, such as minimizing the number of false alarms and generating collision avoidance maneuvers that also ensure separation from the ground.

These goals for the final system behavior create intermediate requirements for each of the four sub-functions in the alerting task:

- *Information Processing*: The current traffic information given to the alerting and decision-making sub-functions must be good enough to enable the desired accuracy of alerts and avoidance maneuver selection. A trade-off in information processing often exists between information accuracy and information latency; for example, the current TCAS II system filters its sensor information to lessen the effect of spurious and erratic sensor readings, and in doing so can have a delay of several seconds in recognizing the deviation of another aircraft. In collision avoidance systems for other phases of flight, this trade-off is not as critical as in parallel approaches, where latency must not be added during the processing of the sensor information.

- *Hazard Assessment and Alerting*: Alerts must be given in time to prevent an impending collision. An alert can be categorized by whether it correctly identifies a potential conflict, and whether it is early enough to allow an avoidance maneuver to be effective. These categories are shown in Table 2.1.

A trade-off exists between the rates of False Alarms and Missed Detections. In order to prevent the catastrophic effects of a loss of aircraft separation, the alerting system needs to be designed conservatively. Conservative design increases the number of False Alarms caused by variability in sensor inputs and uncertainty about the future aircraft trajectories. This requires careful setting of the alerting criteria -- and their implicit threshold -- to generate the best tradeoff between achieving a very low probability of a Missed Detection with as few False Alarms as possible. (Kuchar, 1996)

Table 2.1 Categories of Alerts

Category of Alert	Description
Correct Detection	The alerting system issues an alert which prevents a loss of aircraft separation from occurring.
Correct Rejection	The alerting system does not issue an alert, and the separation between the aircraft without an avoidance maneuver is always adequate.
False Alarm	The alerting system issues an alert in a case where an avoidance maneuver is not required to maintain adequate separation between the aircraft.
Missed Detection	The alerting system does not issue an alert in a case where a loss of aircraft separation occurs.
Late Alert	The alerting system issues an alert too late for an avoidance maneuver to generate adequate aircraft separation.
Induced Collision	The alerting system issues an alert; the resulting avoidance maneuver causes a loss of aircraft separation where the nominal trajectory would have maintained adequate aircraft separation.

- **Decision Making and Resolution Selection:** Given the time to collision when the alert is issued, this sub-function must determine an avoidance maneuver which provides adequate aircraft separation. This decision may involve trade-offs with other concerns, including terrain avoidance and minimal loads on the aircraft due to the maneuver strength.
- **Control Actuation:** Avoidance maneuvers must be executed with as little possible delay, and with a precision that causes the actual escape trajectory to match that desired.

Trade-offs within the sub-functions were already mentioned. A trade-off can also be found between the decision-making, control actuation and alerting sub-functions, described in Table 2.2.

Table 2.2 Trade-Off Between Alerting Accuracy and Avoidance Maneuver Effectiveness

Simple, Early Maneuver	Later, More Sophisticated Maneuver
<p>Always Use the Same, Simple Avoidance Maneuver, Even if it Sometimes Requires an Earlier, Possibly Less-Accurate Alert</p> <p>Pro: Pilot May Be Able to Execute This Maneuver Easier and Quicker (Simpler Decision-Making, Easier & Faster Control Actuation)</p> <p>Con: Some Early Alerts and False Alarms May Be Given (Limited Alerting Performance)</p>	<p>vs.</p> <p>Allow the Alerts to be Delayed for Greater Alert Accuracy by Calculating a 'Tailor Fit' Avoidance Maneuver for Each Situation</p> <p>Pro: Greater Certainty in the Alert Being Correct (Increased Alerting Performance)</p> <p>Con: Pilot May Take Longer to Execute This Maneuver (or Not Agree With It) (More Complex Decision-Making, Slower Control Actuation)</p>

The current TCAS II system resolves this trade-off by always alerting at a specified projected time to collision, and selecting an avoidance maneuver from a limited set of choices. Prototype alerting logic for parallel approaches tend to assume a single standard avoidance maneuver which is easy to understand, which includes a climb to avoid terrain, and which places acceptable loads on the aircraft. (RTCA, 1983; Carpenter & Kuchar, 1996).

At the start of this work, the relative benefits of different roles for an automatic alerting system were not known. Improved methods of displaying the current traffic situation and specialized alerting logic and avoidance maneuver selection logic were suggested; the need for these changes from the current alerting system was uncertain. The next chapter will discuss a preliminary simulator evaluation aimed at these issues.

The task of collision avoidance during closely-spaced parallel approaches has qualities which make it susceptible to pilot non-conformance to automatically displayed alerts and avoidance maneuvers. The potentially non-intuitive nature of the Kill Zone and the many simultaneous criteria an avoidance maneuver must satisfy create many opportunities for disagreement between the pilot and automatic decisions. The time-critical nature of the task leaves little room for these disagreements to be worked out and a safe avoidance maneuver to be executed.

3. Experiment #1: Simulator Evaluation of the Characteristics of a Collision Avoidance System for Parallel Approaches

3.1 Experiment Objectives

At the start of this study, relatively little information was available about the manner in which a pilot would use a cockpit collision avoidance system during closely spaced parallel approaches without intervention from Air Traffic Control. Therefore, a simulator evaluation was conducted, using the MIT Advanced Cockpit Simulator, with the objective of examining the following issues:

- *Ability of Pilots to Avoid Blundering Traffic* At the time of the experiment, no specific data was available indicating whether pilots would be able to successfully avoid blundering traffic given the high workload during final approach, the short time available to initiate the maneuver, and the unfamiliarity of the procedure. Therefore, this experiment was designed to provide a preliminary assessment of the pilots' collision avoidance abilities.
- *Pilots' Use of an Alerting System Displaying Alerts and Avoidance Maneuvers* Because the underlying logic for automatic alerts and avoidance maneuver commands was being prototyped by other studies, this experiment assessed the incremental performance gained when an alerting system provides pilots with executive alerts and avoidance maneuvers.
- *Traffic Display Concepts* Because closely-spaced parallel approaches presents close-range, atypical relative aircraft positions, new displays for illustrating the state of nearby aircraft may be required. This experiment tested several display concepts for this phase of flight.
- *Effects of Flying the Approach on Autopilot* Because this phase of flight is workload intensive, flying the approach on autopilot may allow the pilot to direct more attention to the collision avoidance task. However, should an avoidance maneuver be commanded, more immediate responses may be required than is possible from autopilot-commanded maneuvers or can be achieved after the delay of disconnecting the autopilot. This experiment compared these effects.

The experiment provided data relevant to each of these issues. However, the general results also raised the larger concern of pilot conformance to alerts and avoidance maneuvers commanded by the executive alerting system.

3.2 Experiment Design

3.2.1 Experimental Procedure

Pilots were responsible for intercepting and flying approaches with aircraft on a parallel approach at a runway separation of 4300 feet, the current legal minimum for independent approaches without specialized systems. At all points during the approach, the pilots were told to monitor the other aircraft. In each approach scenario, the other aircraft blundered towards the subject's aircraft at some point during the approach; these blunders could create very hazardous or less hazardous collision situations. The pilots were told to avoid the encroaching traffic by the safest means possible. In all cases, the pilots had a display of the position of the other aircraft. In some cases, the assistance of a TCAS II-type system presented executive traffic alerts and avoidance maneuvers to the pilots; in normal operations, pilots are expected to follow these commands. In the remainder of the cases, no automatic alerts or avoidance maneuvers were displayed. The scenarios were ended when the pilots were clear of the traffic.

An experimenter acted as a second pilot. In some scenarios, the subject-pilots flew the approach manually, and the experimenter acted as the Pilot Not Flying (PNF). In the remaining scenarios the approach was flown by the autopilot with the experimenter acting as the Pilot Flying (PF) until the subject-pilot took manual control to fly the avoidance maneuver. The experimenter never discussed the traffic situation with the subject-pilot or suggested any avoidance maneuver.

Before the measurement runs, the pilots were allowed sufficient training approaches to feel comfortable with the simulator and familiar with the parallel approach scenario. No advice or training on the relative merits of different possible avoidance maneuvers was given. After the experiment runs, the pilots were asked for ratings of the displays and procedures, and for comments on parallel approaches.

3.2.2 Independent Variables

3.2.2.1 Traffic Displays

Several displays of current traffic state information were tested. Each display was designed to examine a specific display concept. All displays assumed the other aircraft's horizontal position and relative altitude information would be available and updated once per second. In addition, each display assumed an onboard database would have knowledge of the position of both the own aircraft's landing runway and the parallel runway; this knowledge is already available in current Flight Management System (FMS) databases.

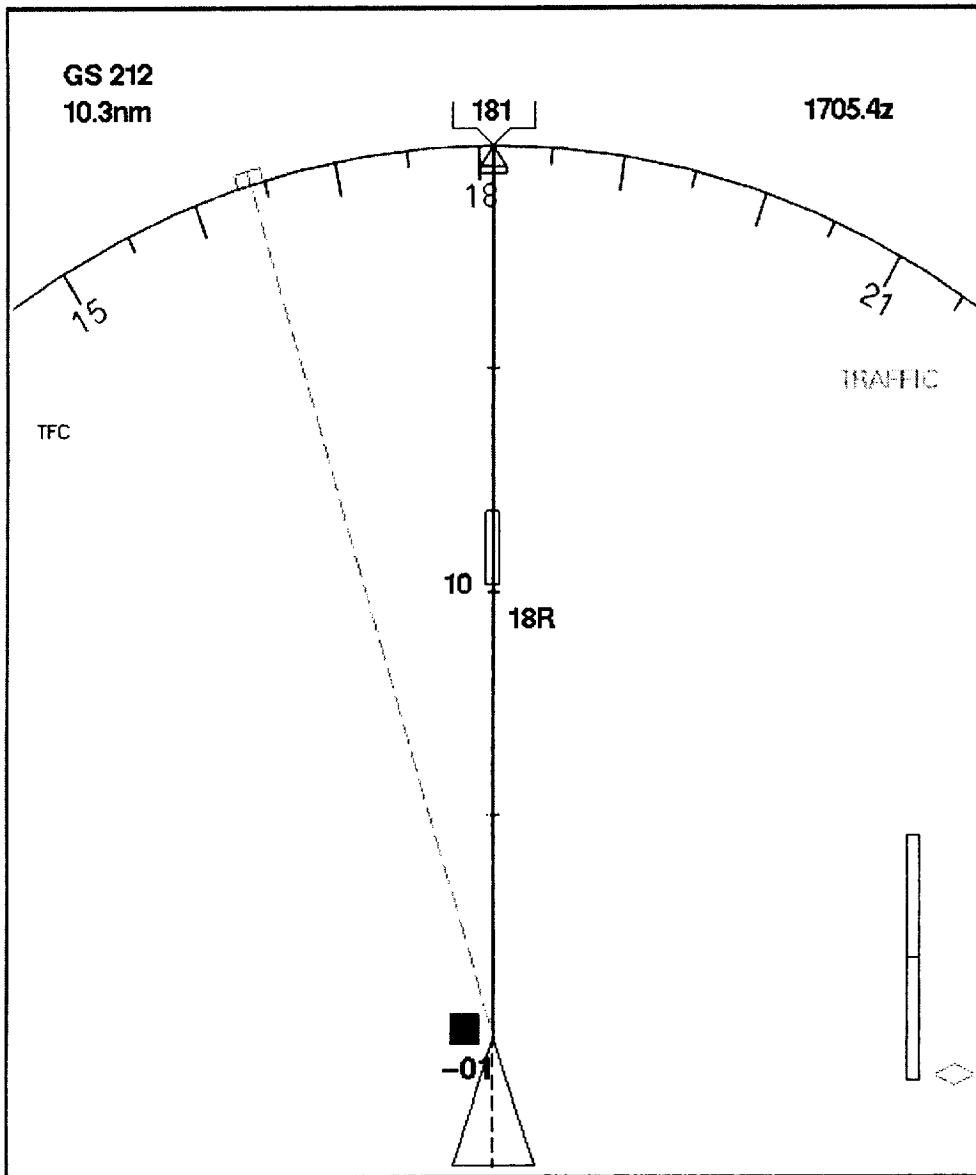


Figure 3.1 Baseline Traffic Display, Showing Traffic on Parallel Approach (Grayscale, Black-White Inverted for Clarity, 90% Scale)

Baseline EHSI As shown in Figure 3.1, the Baseline EHSI display uses the same format as a current implementation of the TCAS II traffic display, which is based on the moving map or Electronic Horizontal Situation Indicator (EHSI). This display presents navigation information to the pilot in the horizontal plane from a top-down view. The own aircraft position is indicated by the fixed triangle at the bottom center of the screen. The position of navigation features are drawn relative to the own aircraft symbol. For example, in Figure 3.1 the landing runway, labeled '18R', is shown 10 miles in front of the own aircraft. As the subject continues the approach, the runway will appear to move down the screen.

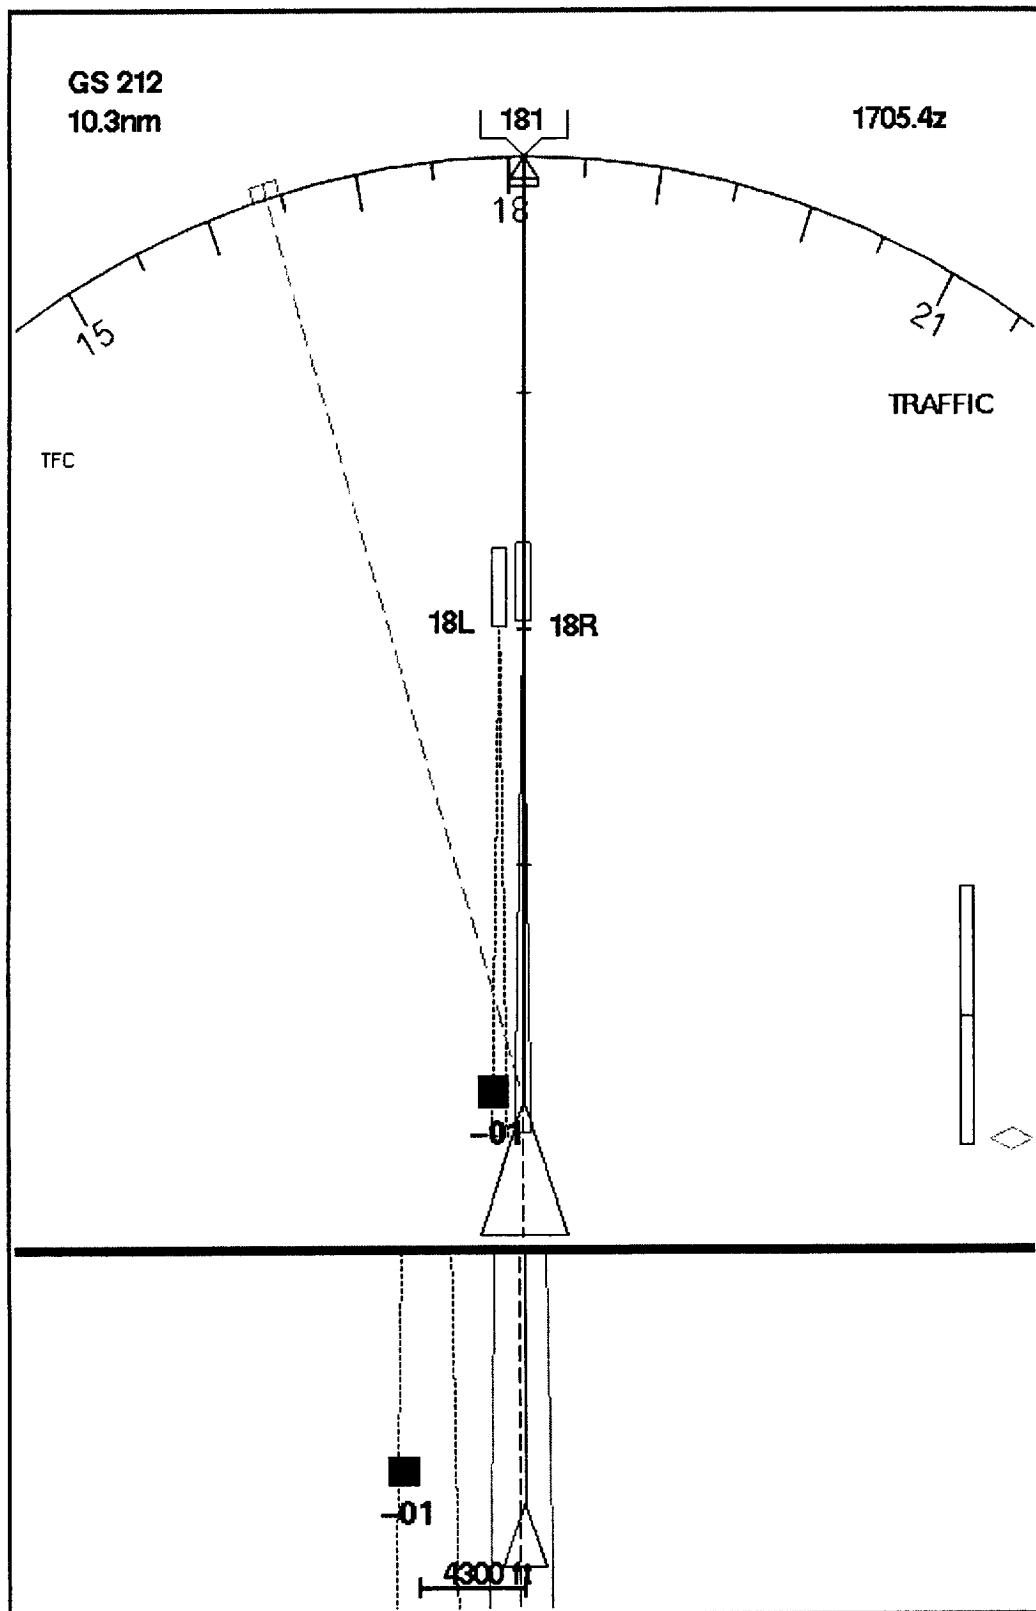
The other aircraft's position is drawn relative to the own aircraft symbol. The relative altitude, discretized to hundreds of feet, is shown in text. The relative position of the text, above or below the traffic symbol, provides an additional indicator of relative altitude. When the aircraft climbs or descends beyond a threshold rate, a vertical trend arrow appears next to the aircraft's symbol.

The shape of the traffic symbol follows the convention used by the TCAS II system. Normally a hollow white diamond, it becomes a filled white diamond when the traffic is within 6 miles horizontally and 3000 feet vertically of the own aircraft; during parallel approaches, the traffic almost always falls within this criteria. When the TCAS II system generates an warning Traffic Alert (TA), the symbol becomes a yellow circle. When the TCAS II system generates an executive Resolution Advisory (RA), the symbol becomes a red square. In the experiment runs where the TCAS II system did not present alerts and avoidance maneuvers, the symbol remained a white diamond.

Ahead of the own aircraft symbol is an indication of the approach path and the landing runway. No information is shown about the parallel approach or runway. In Figure 3.1 the display is shown at its smallest normal range, 10 miles. In this experiment, however, pilots were allowed to select a smaller range if they desired better resolution of the relative position of the traffic; the selection of a smaller range could scale the navigation information, such as the landing runway, off the screen.

Enhanced EHSI As shown in Figure 3.2, this display is based upon the TCAS display, with two new features. First, information is shown about both the ownship's approach path and the parallel approach path, allowing for easy comparison of the other aircraft's position to its approach path. In addition to the own aircraft's landing runway, the position of the own aircraft's approach path is represented in magenta by an 'approach fan'. Its shape correlates to the fan-shaped localizer beam used by the Instrument Landing System (ILS). The width of the 'approach fan' correlates to the distance shown on the cockpit landing guidance displays as one 'dot' from center. The parallel runway and its approach fan are shown in white.

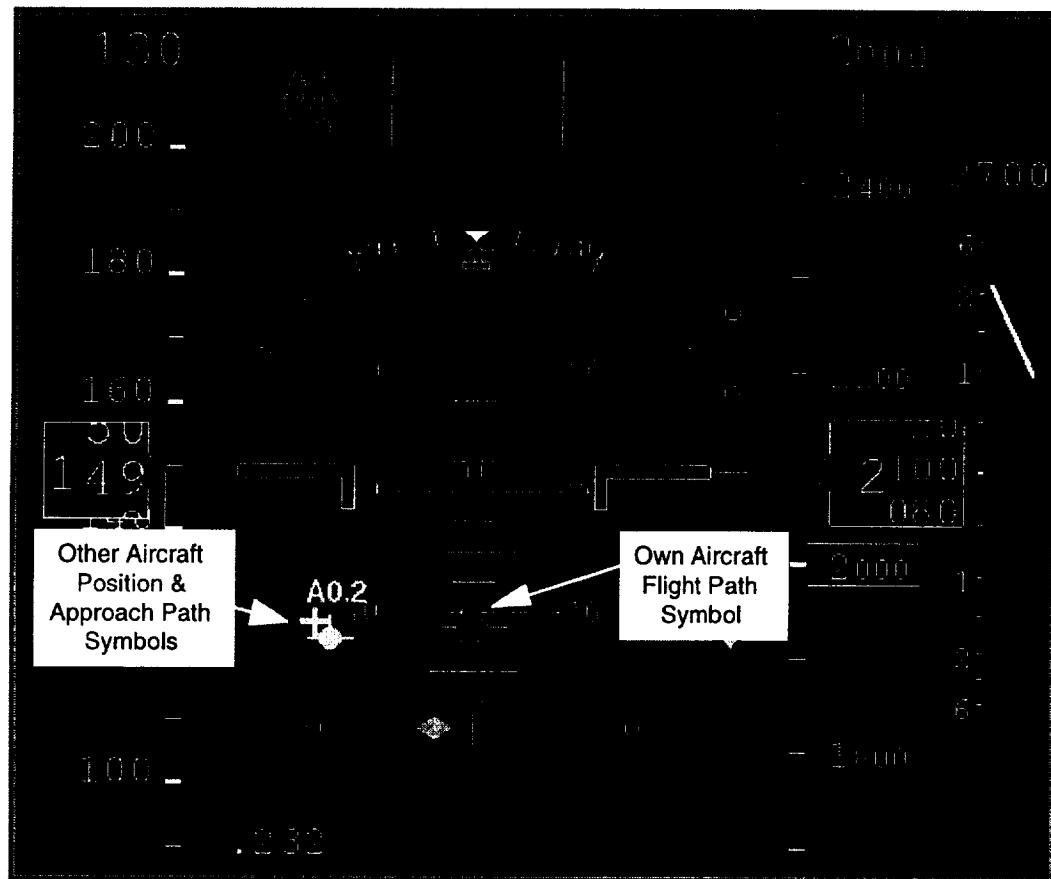
Second, because different display resolutions are needed during this phase of flight to scan for the far-away navigation information and the close-in traffic information, a secondary display was added at the bottom. This display has a display resolution four times that selected for the EHSI, and scales with the EHSI's range setting. With this secondary display, it was intended pilots could find desired navigation and traffic information without requiring a change in the display range.



**Figure 3.2. Enhanced EHSI Traffic Display
(Grayscale, Black-White Inverted for Clarity, 90% Scale)**

Traffic Display on PFD The next two display combinations examined the presentation of traffic information on the pilots' Primary Flight Display (PFD), as shown in Figure 3.3. The PFD provides the state information required to control the own aircraft during approach, including aircraft attitude, localizer and glideslope deviations, airspeed, altitude, vertical speed, and autopilot mode and target states. Normally flight director bars can also be selected; this simulator did not have this capability.

The traffic information is drawn in a vertical plane with a viewpoint looking along the approach path. Using this viewpoint, the altitude of the parallel aircraft is drawn, relative to the own aircraft, as the height on the display of the aircraft's symbol relative to the flight path angle symbol. The across-track position of the parallel aircraft is drawn, relative to the own aircraft, as the side-to-side distance on the display. This viewpoint does not provide a graphical indication of the along-



**Figure 3.3. Traffic Display on PFD
(80% Scale, Gray Scale)**

track spacing of the aircraft. Therefore, the distance the other aircraft is ahead or behind the own aircraft is shown in text in nautical miles. If the parallel aircraft is ahead, the text is prefixed with an 'A'; if the parallel aircraft is behind the own aircraft, the text is prefixed with a 'B'. Cross-hairs are drawn to indicate the one-dot localizer and glideslope position of the other aircraft.

Altogether, four display combinations were tested. The first two showed the baseline EHSI display alone and the enhanced EHSI display alone. The third display combination presented the traffic display on the PFD in conjunction with the baseline EHSI display; this combination tested the ability of the traffic information on the PFD to provide a principal source of traffic information, as the baseline EHSI display provided less traffic information. The fourth display combination presented the traffic display on the PFD in conjunction with the enhanced EHSI display.

3.2.2.2 Alerting System and Autopilot Availability

To examine the comparative benefit of having the alerting system display executive-type alerts and avoidance maneuvers for the pilot to follow, the TCAS II alerts and display of avoidance maneuvers were turned off in some experiment runs. These results were desired to determine the need for alerting and decision-making systems; if required, these systems would have to be developed specifically for the parallel approach geometry.

In addition, approaches where the subject acted as Pilot Flying (PF) and flew the approach manually were compared to approaches where the subject acted as Pilot Not Flying (PNF) and the approach was flown on autopilot. This comparison examined the effects of workload on performance and the merits of flying the approaches on autopilot.

These different conditions were described to the subject-pilots as different 'procedures'. Only three combinations of these two variables were tested, in order to keep the length of the experiment manageable for the pilots. They were:

- Autopilot approach, subject PNF, TCAS alerts and avoidance maneuvers displayed.
- Manual approach, subject PF, TCAS alerts and avoidance maneuvers displayed.
- Manual approach, subject PF, no TCAS alerts or avoidance maneuvers displayed.

3.2.2.3 Scenarios

Three approaches were flown under each display-procedure condition. Each involved a different type of traffic scenario. The scenarios were set to occur at different times in different

runs, reducing predictability. The intruding aircraft was presumed to be installed with TCAS II, and in each case followed the maneuver it commanded.

In the 'Missed Intercept' scenario, the other aircraft never intercepted the parallel approach path, but instead continued towards the subject. If neither the subject nor the intruder flew an avoidance maneuver, they would pass within 500 feet. Because the intruding aircraft did not require time to establish a collision course, this scenario created the most time-critical hazard.

In the 'Hazardous Blunder' scenario, the parallel aircraft intercepted and flew its own approach until a pre-selected point, where it turned towards the subject. Without an avoidance maneuver, the two aircraft would pass within 500 feet of each other. The collision, starting from when the intruder started to deviate, was specified to occur within 30 to 45 seconds.

The 'Less-Hazardous Blunder' scenario was similar to the 'Hazardous Blunder' scenario, except the intruding aircraft was steered to pass approximately 1500 feet away from the subject. These scenarios also generated TCAS II alerts and avoidance maneuvers.

3.2.2.4 Test Matrix

The test matrix was three dimensional, varying the type of traffic displays, the approach scenarios, and various combinations of autopilot control and automatic alerts, as shown in Figure 3.4. Each subject flew a total of 36 experiment runs.

Alerting System & Autopilot Availability	Displays ->			
	Baseline EHSI	Enhanced EHSI	Combined Baseline EHSI & PFD	Combined Enhanced EHSI & PFD
Autopilot Approach (Before Alert) Follow TCAS II Maneuvers	3 Scenarios	3 Scenarios	3 Scenarios	3 Scenarios
Manual Control (Before Alert) Follow TCAS II Maneuvers	3 Scenarios	3 Scenarios	3 Scenarios	3 Scenarios
Autopilot Approach (Before Action) Pilot Decides on Action & Maneuvers	3 Scenarios	3 Scenarios	3 Scenarios	3 Scenarios

Figure 3.4 Experiment #1 Test Matrix

3.2.3 Simulator Setup

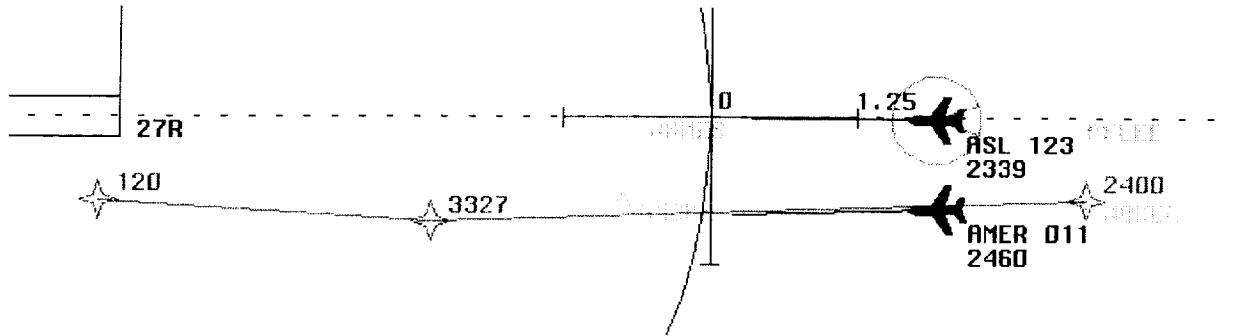
The MIT Advanced Cockpit Simulator was based upon a Silicon Graphics workstation. The computer provided both the graphics emulating the cockpit displays and the computation to simulate the aircraft dynamics and to drive the ancillary controls. The simulated aircraft had a level of performance approximating a Boeing 737.

Several methods were available to the subjects for controlling the aircraft. A side-stick was provided for manual control. A Mode Control Panel (MCP) was provided; due to the experimental procedures, this was only used during autopilot approaches and for adjusting the speed commanded by the autothrottles. Flaps were commanded by a lever on the throttle quadrant; gear and display features were commanded through keyboard inputs by an experimenter acting as Pilot Not Flying (PNF). A Control Display Unit (CDU) and throttles were present, but not needed.

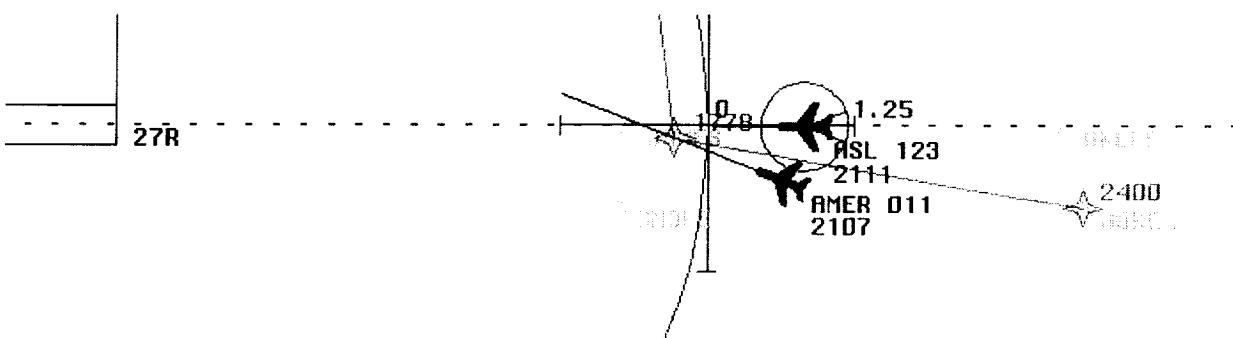
The Primary Flight Display (PFD) and Electronic Horizontal Situation Indicator (EHSI) were shown on the workstation's display, directly in front of the pilot. The traffic information was integrated onto these displays. Ancillary displays were also shown, including gear and flap position indicators, and a display of the autopilot's settings. Because the approaches were in simulated Instrument Meteorological Conditions, no out-the-window view was given.

Traffic alerts and avoidance maneuvers were based on an early version of the TCAS II logic. A slight modification fixed the alerting sensitivity so that consistent alerts would be given down to the runway threshold; the system was prevented from commanding a descent into the ground.

The dynamic states of the parallel aircraft were updated by a Robust Situation Generation system, running on a separate machine. (Johnson & Hansman, 1995) This system included a Simulator Control Station for scripting scenarios before the experiment and for monitoring the aircraft position during experiment runs. During experiment runs this system steered the aircraft according to criteria in the pre-determined scripts. These criteria could be determined relative to the subject position, allowing the desired traffic situation to occur reliably despite the subject's independent decisions to change the own aircraft's speed and flap settings during the approach.



a) Flight 011 Appears to be Flying a Normal Approach to Runway 27L
Commanded Waypoints Accelerating It Forward Into 'Kill Position'



b) Flight 011 Blundering Towards Subject (ASL 123)
Once it Reached 'Kill Position', New Collision Trajectory Commanded Using Subject
Relative Waypoints

Figure 3.5 Simulator Control Station at Two Points During a 'Hazardous Blunder' Scenario

An example scenario is shown in Figure 3.5. In the first snapshot from the simulator control station screen, the intruder, Flight 011, is being commanded to fly slightly faster than the subject aircraft, Flight 123. Once the aircraft are in the relative position where a turn towards the subject causes a hazardous condition, the intruder is automatically given a new commanded flightpath which commands the intruder to home in on the subject. This homing continues until the intruder receives a TCAS resolution advisory (RA), which commands an avoidance maneuver. The intruder follows this maneuver for the duration of the alert. Once the alert is over, the intruder steers to a missed approach point away from the airport.

3.2.4 Subjects

Eighteen airline pilots participated in the experiment from two major airlines. All but one were current in glass cockpit aircraft. Ten were captains, the remaining eight were first officers. The subjects had an average of 15,200 total flight hours. They reported having previously experienced an average of seven TCAS Resolution Advisories (RAs) in actual operations.

3.2.5 Measurements and Performance Metrics

Several measurements were taken, including:

- Complete state data of the subject's aircraft and the intruding aircraft in each run.
- Time when the pilot initiated an avoidance maneuver in each run.
- Time when the TCAS alerts were generated in each run.
- Subjective comments of the pilots at the completion of the experiment.

From these measurements, several performance metrics were calculated:

- Whether the pilot's decision to initiate an avoidance maneuver was a correct detection of an impending collision hazard.
- The type of avoidance maneuver, measured by the aircraft bank and pitch rate.
- The minimum miss distance achieved between the subject and intruding aircraft.
- Subjective ratings of the different experimental conditions.

Several non-exclusive terms will be used to describe these metrics in the next section.

- An 'Early Avoidance Maneuver' was an avoidance maneuver executed by the pilot while the other aircraft was still following its nominal approach and before the TCAS provided (or would have provided, when the alerts were not shown) a commanded maneuver.
- A 'Turning Avoidance Maneuver' is defined as an avoidance maneuver in which the aircraft bank exceeded five degrees from level. Distinctions were also made between 'Turn Away' (from the intruder), 'Turn Into' (the intruder), and 'Turn Both Ways' maneuvers.
- A 'Near Miss' is defined as an event where the miss distance between the subject and intruding aircraft was less than 500 feet, measured between centers of mass. The number of approaches which resulted in a 'Near Miss' was tallied. Because this 500 feet criterion may be considered an absolute minimum value, the number of approaches which resulted in a miss distance under 1000 feet will also be given for comparison.

Differences between measurements from different conditions were tested for statistical significance using unpaired t-tests. Levels of statistical significance will be noted in the results. (e.g. Rice, 1988; Hogg & Ledolter, 1993)

3.3 Experimental Results

This section details the results of this experiment. First, the conformance of the pilots to automatically displayed alerts and its correlation with traffic displays and scenario type are discussed. Then, the effects of each element of the test matrix -- display, alerting system availability and scenario type -- are examined.

3.3.1 Conformance of Pilots to the Automatic Avoidance Maneuvers

In two of the three procedures tested, an alerting system, based on the current TCAS II system, presented pilots with executive alerts and avoidance maneuvers to follow. At the completion of the experiment, the pilots' trajectories were compared with the avoidance maneuvers commanded by TCAS to ascertain whether the pilots had conformed to the alerting system.

The TCAS logic for generating a vertical avoidance maneuver was designed assuming a five second pilot reaction time, followed by a 0.25g pull-up or push-over maneuver to the commanded pitch attitude. In analyzing the results of this experiment, any maneuver that met or exceeded this criterion was considered to be in conformance with the TCAS command when it was shown, or was considered to match the TCAS command when the pilot was not shown the TCAS maneuver. This is illustrated in Figure 3.6 for a commanded climb; the criteria for a commanded descent

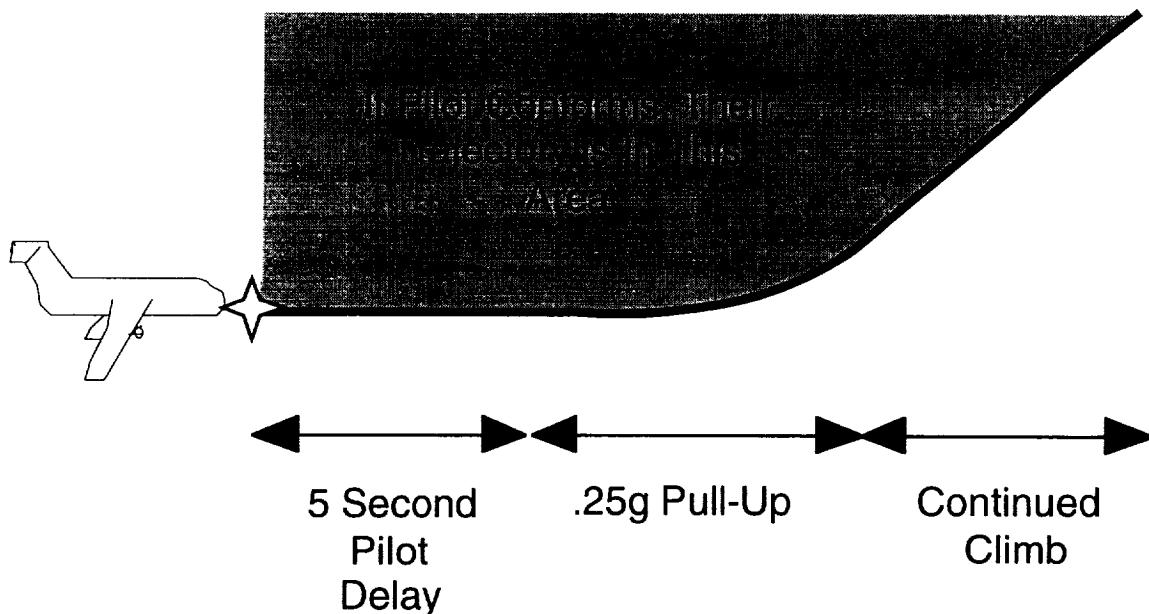
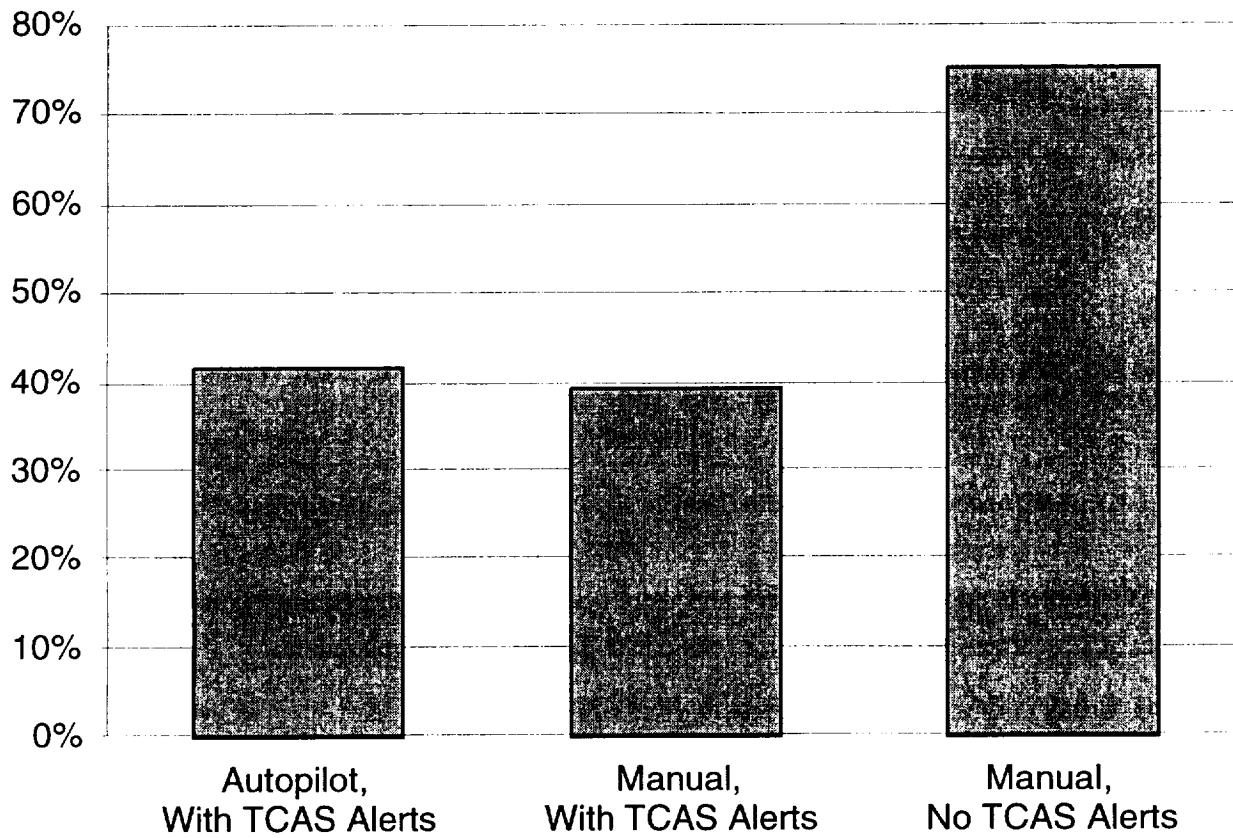


Figure 3.6 Criteria for Judging Conformance to a TCAS Selected Avoidance Maneuver

followed the symmetric, but opposite, trend. Whether the pilot's avoidance maneuver also included a turn component was not considered in this definition.

The display of the TCAS generated maneuver had a significant impact on whether the pilot's maneuver met the conformance criterion. As shown in Figure 3.7, the pilots maneuver did not match what the TCAS would have commanded in slightly more than 75% of the approaches where the automatic information was not displayed, indicating pilots tended to choose different avoidance maneuvers when an alerting system was not available. When the automatic information was displayed, it was not conformed to in approximately 40% of the approaches, with no significant difference in frequency between approaches flown manually and on autopilot.



**Figure 3.7 Frequency of Pilots' Maneuvers Not Conforming to (or Not Matching) the TCAS Suggested Maneuver, by Procedure
(n is approximately 190 in each condition)**

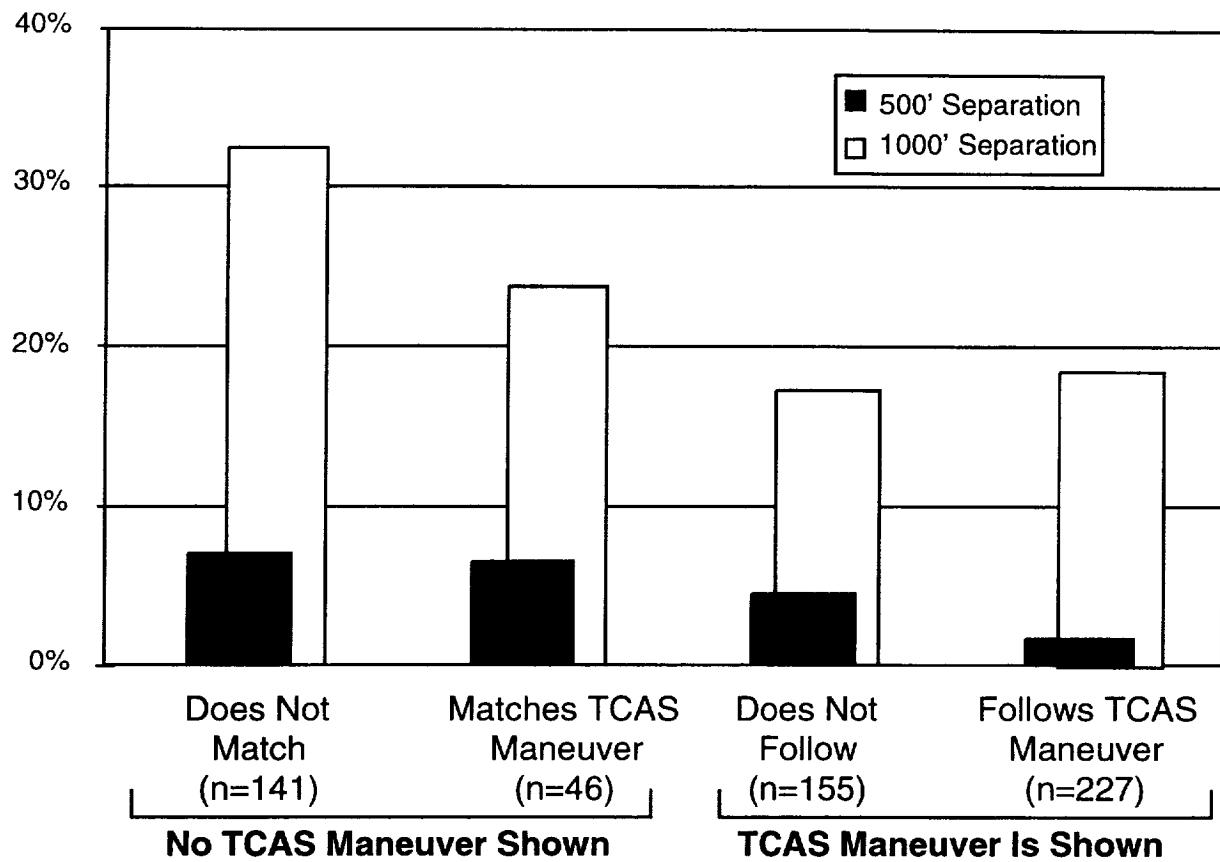


Figure 3.8 Frequency of Near-Misses in Conditions With and Without Display of Avoidance Maneuvers, and With and Without Pilot Conformance to the Maneuver

This non-conformance was found to have a significant impact on performance as measured by the frequency of near-misses. Figure 3.8 shows the frequency of near-misses for the conditions where TCAS alerts and maneuver were and were not shown, and piloted maneuvers did and did not meet at least the minimum specifications of the TCAS selected maneuver.

Two results can be noted here. First, a trend of fewer near-misses is visible both with the availability of TCAS alerts & commanded avoidance maneuvers, and with pilot conformance to them. This drop is statistically significant between the extreme cases where TCAS commands were not shown and the pilots did not match the TCAS maneuver, compared to cases where the pilots were shown the TCAS maneuver and conformed to it ($p < 0.05$).

Second, a significant increase in collision avoidance performance (with the 1000 foot separation criteria) is found in the cases where the TCAS maneuver is displayed, even if it is not exactly followed ($p < 0.01$), suggesting the alerting function of the automatic alerting system provided an additional benefit.

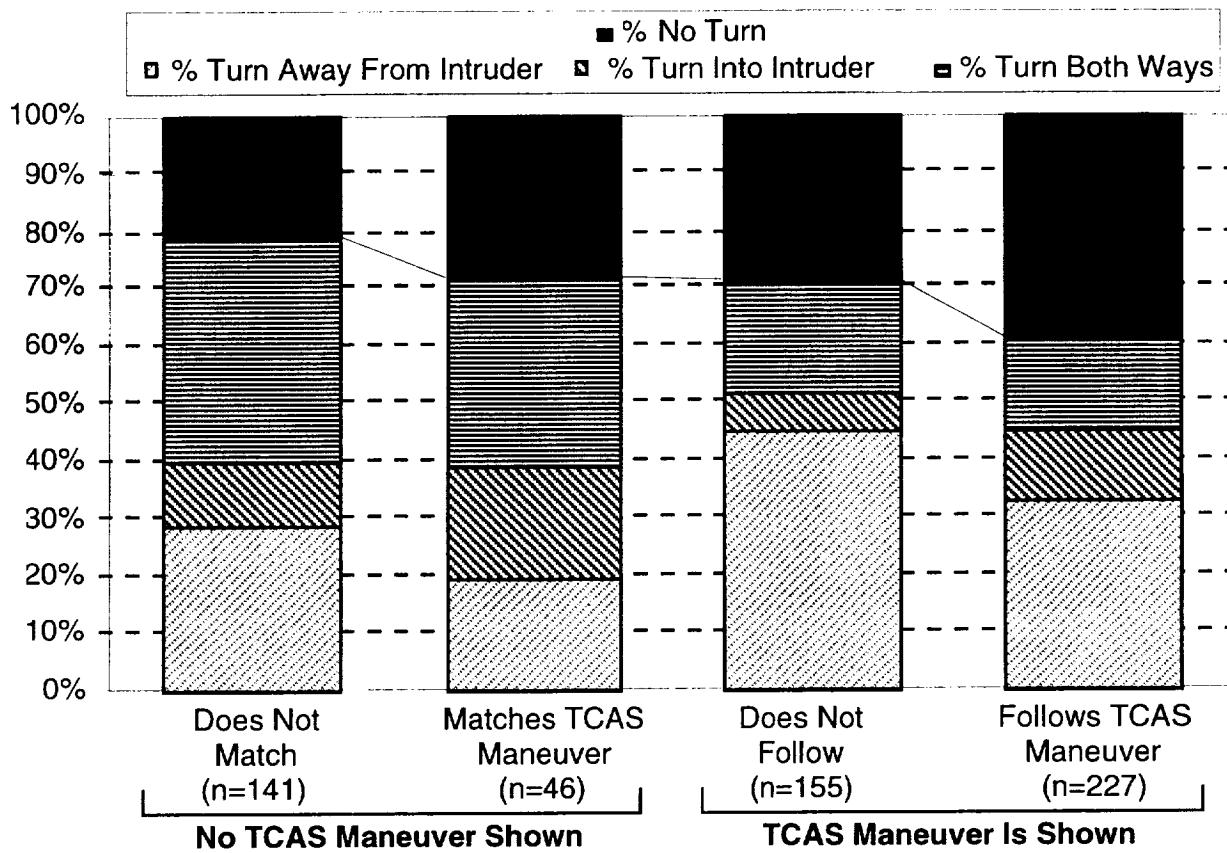


Figure 3.9 Frequency of Each Type of Turning Maneuver, Compared with the Display of, and Conformance to, TCAS Avoidance Maneuvers

No single causal factor of the high non-conformance rate can be isolated. Pilot reaction time alone does not show a strong effect. 66% of the pilots reacted within the five second allowance assumed by the TCAS system, and of these only 61% matched the displayed TCAS maneuver. Of the pilots who acted shortly before the alert or after the five second allowance (13% and 20% respectively), the trajectory followed by the pilots still frequently matched what the TCAS guidance commanded (71% and 33% respectively).

However, two factors appear to be correlated with pilot non-conformance to displayed TCAS maneuvers. First, non-conformance to the (vertical) TCAS maneuver may be related to the turning maneuvers that the pilots often performed at the same time, as shown in Figure 3.9. Overall, pilots did not turn in 32% of the approaches (i.e. the maximum bank angle after the alert was less than five degrees); 34% of the time the pilots turned away from the intruder, 11% of the time pilots turned toward the intruder, and 23% of the time pilots turned one way and then another. Pilots who followed a displayed TCAS maneuver turned significantly less often than pilots who did not follow the TCAS maneuver ($p < 0.05$), as shown in Figure 3.9. This may suggest that the pilots, by executing a turn, felt the commanded vertical maneuver was no longer required.

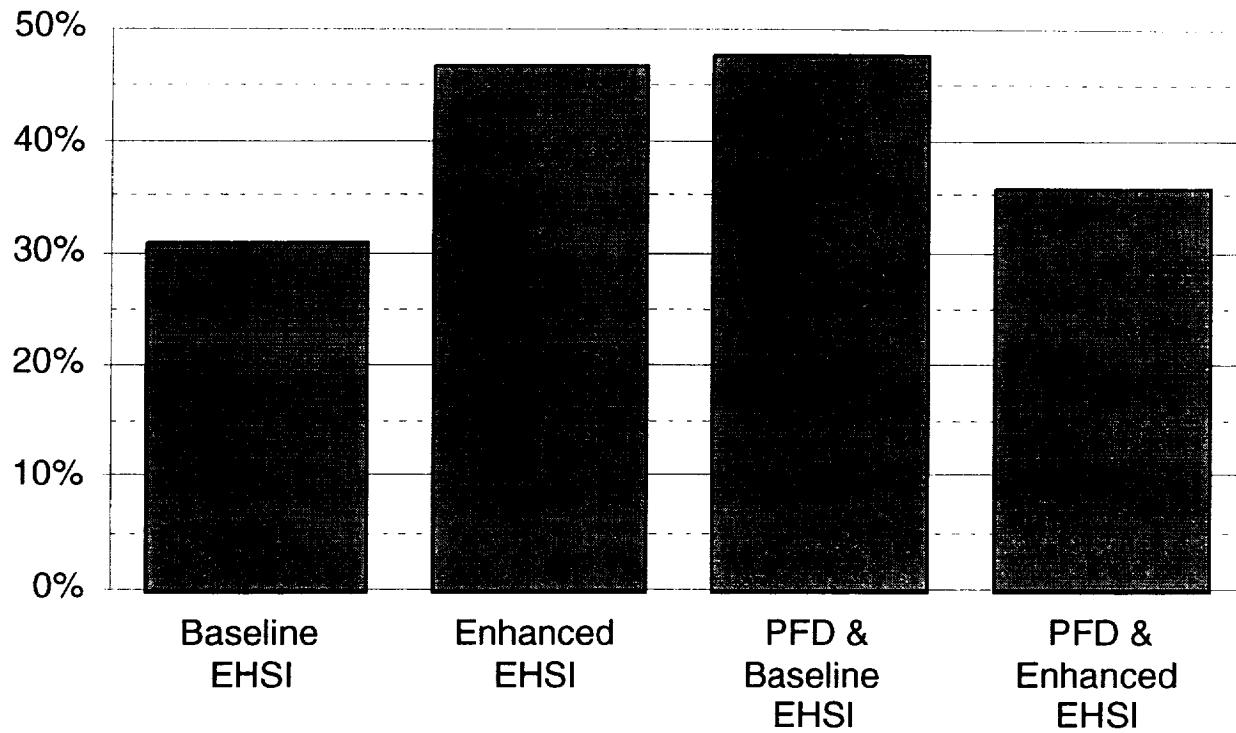


Figure 3.10 Frequency Pilot's Maneuver Matched TCAS Maneuver, by Traffic Display (n= 130-150 in each display condition)

Second, pilot non-conformance to commanded avoidance maneuvers was found to increase significantly with the enhanced displays, as shown in Figure 3.10. The difference between the non-conformance rate with the baseline display and with the enhanced EHSI display is very significant ($p < 0.01$); the difference between the baseline and Primary Flight Display (PFD) is significant ($p < 0.05$). This higher non-conformance rate with the prototype displays was found to correspond with a higher rate of near-misses with the prototype displays.

The non-conformance rate was also found to vary with the type of scenario. The non-conformance rate was significantly higher for the 'Hazardous Blunder' type scenarios than the 'Less Hazardous' type scenarios ($p < 0.05$) and the 'Missed Intercept' type scenarios ($p < 0.01$). This may indicate that pilots had greater confidence in the alerts and avoidance maneuvers generated for the hazardous blunder scenarios than for the other scenarios, or that they chose to react differently to this scenario, in a maneuver more consistent with that commanded by the alerting system.

Early Avoidance Maneuvers were executed in 17% of the approaches. This figure, while significant, may indicate any of three factors. First, given the experimental procedure, pilots may

have expected a blunder to occur, causing their reactions to be unnaturally conservative. Second, the traffic displays used in the experiment may have over-emphasized the threat posed by aircraft on a parallel approach. However, this high false alarm rate may also indicate a lack of confidence by pilots in the safety of parallel approaches and in the effectiveness of TCAS alerts; as such, these Early Avoidance Maneuvers may also represent instances of non-conformance, where the pilots were not waiting for the automatic alerts but were instead electing to initiate an avoidance maneuver based on other criteria.

3.3.2 Effects of Traffic Displays

The previous section discussed the pilot's higher non-conformance rate with the prototype displays and their corresponding higher rate of Near Misses. This section will discuss the possible characteristics of the prototype displays which may have contributed to a lower conformance rate.

Subjective ratings indicate that the pilots preferred the display combinations which had either or both of the two enhanced displays: the Enhanced EHSI display and the display of traffic on the Primary Flight Display (PFD). In order to generate quantifiable subjective ratings, pilots were asked at the completion of the experiment to give paired-comparison ratings between each of the displays. The results were analyzed using the Analytic Hierarchy Process (AHP). (Method described in Yang & Hansman, 1995) These AHP ratings were normalized to sum to one, and their relative percentages are shown in Figure 3.11.

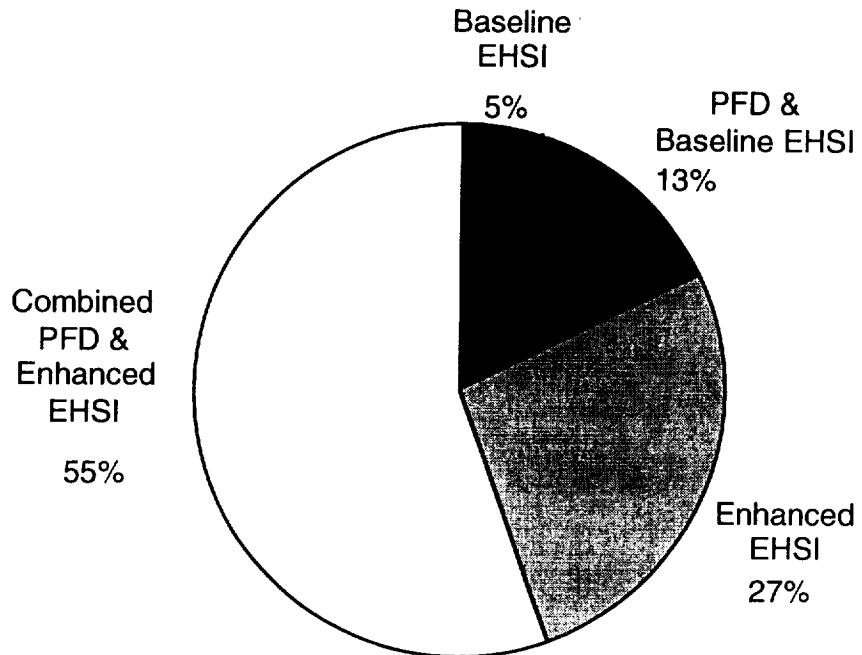


Figure 3.11. Subjective AHP Ratings of the Displays in the Baseline Experiment for the Question “Which Display Gives A Clearer Picture of the Traffic Situation?” (18 pilots)

The pilots’ comparative approval of the displays can be found by finding the relative magnitude of their ratings. Having the combination of the Enhanced EHSI and PFD displays was considered 11 times better than the TCAS II type baseline display (55% / 5%). The enhanced EHSI alone was considered approximately five times better than the baseline TCAS II display (27% / 5%), and approximately twice as good as the PFD and baseline display combination (27% / 13%). Finally, the combination of the PFD display and the baseline was almost three times better than the baseline display alone (13% / 5%). Overall, this indicates a strong preference for the enhanced displays.

Pilots were invited to give free responses to questions about the features of the traffic displays which they liked. Seventeen of the 18 pilots cited the indications of the other aircraft’s approach path as useful for monitoring the other aircraft’s position relative to its approach path. However, decisions to initiate an avoidance maneuver based on this position indication alone could conflict with TCAS alerts, which also consider convergence rate.

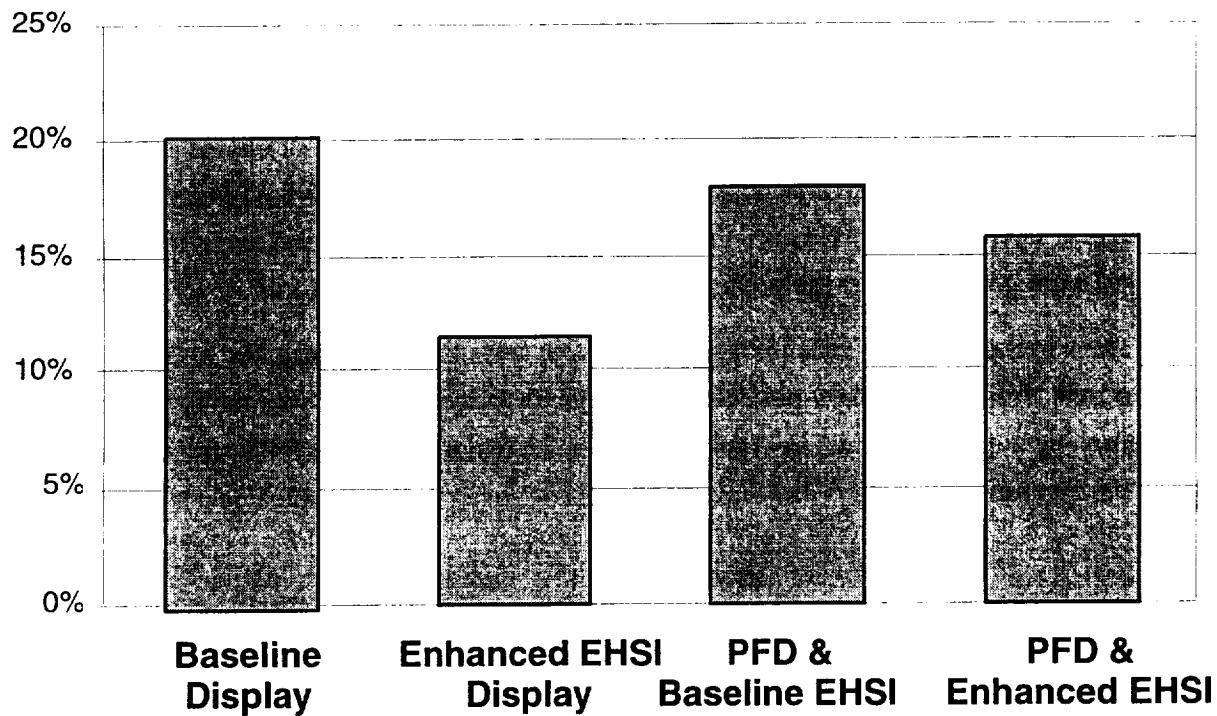


Figure 3.12. Percentage of Approaches Resulting in an Early Avoidance Maneuver, by Traffic Display (n=160 in each display condition)

Pilots executed fewer Early Avoidance Maneuvers when the new traffic displays were available. As shown in Figure 3.12, the drop is very significant ($p < 0.01$) from the baseline display to the enhanced EHSI display. This may be an indication the displays increased pilot confidence in parallel approaches.

Based on these results, pilots appeared to prefer the enhanced traffic displays, and to feel more confident in the parallel approach task when these displays were available. These results may indicate the pilots may have perceived less of a dependency on the automatically generated alerts and avoidance maneuvers when the prototype displays were available.

3.3.3 Effects of Alerting System Availability

A significant benefit in performance was identified, in Section 3.3.1, when automatically generated alerts and avoidance maneuvers were displayed. This section discusses a correlating preference by the pilots for an alerting system. For example, in subjective AHP ratings of the ‘procedures’ used in this experiment, the procedures with TCAS alerts and avoidance maneuvers displayed were preferred over the procedure without them by a factor of four.

To more specifically examine pilot preferences for automatic assistance, the pilots were each asked to rank, from 1 to 8, eight possible roles of an alerting system. These rankings, each represented by a dot, are shown in Figure 3.13. The boxes represent the 94% confidence interval for the median ranking given to each role of an alerting system. Pilots indicated a significant

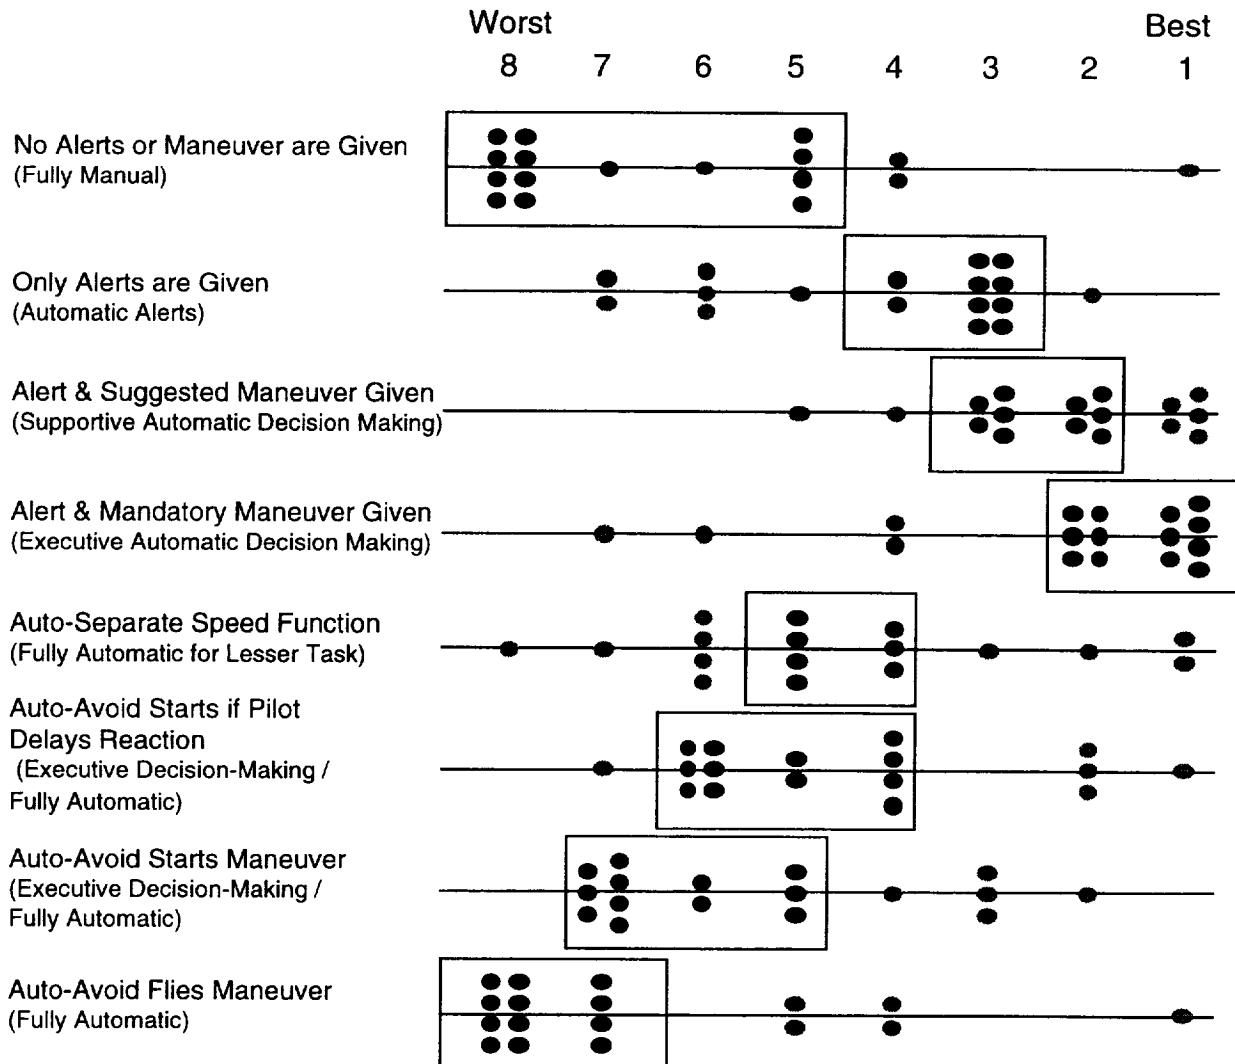


Figure 3.13 Pilot Rankings of Automatic Functions

preference for an alerting system role which displayed alerts and mandatory or suggested avoidance maneuvers, the alerting system role currently provided by TCAS. Fully automatic and fully manual levels of automation were equally disliked and ranked the lowest.

3.4 Summary of Results and Discussion

This experiment tested the ability of pilots to avoid aircraft straying from a parallel approach. The experiment was intended to test several traffic display concepts, the need for automatically generated alerts and avoidance maneuvers, and the relative benefits of flying the approach on autopilot or manually. Several significant effects were identified which have implications for closely spaced parallel approaches.

- Pilots indicated a strong preference for automatically generated alerts and the display of avoidance maneuvers for the pilots to execute. The presence of automatically generated alerts and avoidance maneuvers significantly improved collision avoidance performance.
- The rate of pilot non-conformance to automatically generated alerts and avoidance maneuvers was approximately 40%. This differs from the assumption of conformance implicit in the alerting system's calculation of avoidance maneuvers.
- While a need for an automatic alerting and collision avoidance system for parallel approaches was identified, the performance of the baseline collision avoidance system used in this experiment was not achieved when its commanded avoidance maneuver was not followed. Avoidance maneuvers which did not conform to those generated by the TCAS system had significantly more Near Misses, suggesting pilot conformance to the alerting system is a design issue which impacts total system performance.
- Several factors appeared to be correlated with low pilot conformance to automatic avoidance maneuvers. For example, the pilots' maneuvers included turns significantly more often when they did not conform, suggesting they preferred a turning maneuver over the commanded vertical maneuver.
- Pilot conformance appears to be related to the traffic display available to the pilot during the approach. Two prototype traffic displays tested several display concepts. An Enhanced EHSI display provided an indication of the other aircraft's approach course, and a secondary display gave an expanded scale viewpoint of the traffic immediately around the own aircraft. Traffic was also drawn on the Primary Flight Display (PFD) situating traffic information on the display central to the pilot's scan during approach. The prototype displays produced mixed results. Compared with the baseline current TCAS II traffic display, pilots had a strong preference for the prototype

displays tested in this experiment. However, the pilots' non-conformance rate was significantly higher when the prototype displays were available; this effect is correlated with a lower performance with these displays.

The discrepancies between the pilots' actions and those given by the executive-type alerting system suggest the pilots assessed the collision hazard and possible avoidance maneuvers in parallel with the alerting system, indicating a potential lack of confidence in the commands given by TCAS.

Mismatches between the decisions made by the pilot and the automatic commands given by the TCAS II system may have existed. For example, the enhanced displays provided pilots with an 'approach fan' indicating the cross-track position of a normal approach. The pilots indicated they liked this feature; some commented that it freed them from monitoring the convergence rate of the other aircraft. Therefore, this feature may have unintentionally encouraged an alerting logic based primarily on the other aircraft's lateral deviation from their approach path, and have therefore encouraged a higher non-conformance rate.

Likewise, the displays emphasized horizontal separation from the intruding aircraft, possibly promoting horizontal avoidance maneuvers over the vertical maneuvers commanded by the alerting system. The strategies used by the pilots tended to have lower collision avoidance performance than those used by TCAS II. These results present the problem of developing an alerting system which uses sophisticated, higher performance strategies while encouraging pilot conformance. The next two chapters detail further simulator evaluations of this issue. First, Chapter 4 discusses a simulator evaluation of the strategies preferred by subjects without the presence of an alerting system, in order to understand the type of strategies which subjects are willing to use. Chapter 5 discusses the design and simulator evaluation of an alerting system which attempts to alleviate pilot non-conformance.

4. Experiment #2: Simulator Evaluation of Alerting and Avoidance Strategies Used by Subjects

4.1 *Experiment Objectives*

Experiment #1 identified situations where pilots did not conform to automatically generated alerts and avoidance maneuvers, with a resulting increase in the rate of loss of aircraft separation. Because the traffic displays emphasized features which suggest different alerting and decision-making criteria, it was hypothesized the pilots were performing their own different alerting and avoidance maneuver selection based on the information emphasized in the displays, in parallel with the alerting system. The resulting mismatch between the pilots' and alerting system's decisions may have contributed to the frequent occurrences of pilot non-conformance.

This chapter details Experiment #2, which had the following three objectives:

- Evaluate which criteria appear to be used by subjects to react and decide on an avoidance maneuver in the absence of automatic alerts. Satisfying this objective would allow for identification of discrepancies between the criteria used by the alerting system and by the pilot.
- Ascertain how display features affect a user's ability to detect a conflict. The displays used in Experiment #1 emphasized position. With datalink between aircraft, trend information such as aircraft heading and turn rate may be available. It was hypothesized that showing this trend information would shape the subjects' decision criteria to incorporate convergence rate and encourage a strategy with higher performance.
- Test the effect of subject workload on the subject's ability to detect a potential conflict. It was hypothesized that subjects would revert to simpler decision strategies during periods of high workload.

4.2 Experiment Design

4.2.1 Experimental Procedure

Using a workstation based, part-task simulator, each subject flew several experiment runs. Each run consisted of three sequential parts:

- *The Flight* The subjects were told they were flying an approach. Their primary task was to keep their wings level despite turbulence, referenced to an artificial horizon, through the use of a sidestick. The sidestick commands did not affect the path of the aircraft so that consistent approach paths were followed. Their secondary task was to press a red button on the sidestick as soon as they thought the aircraft on a parallel approach was blundering towards them, as evidenced by the traffic display. The approach paths were separated by 2000 feet.
- *The Maneuver Selection* Once the subjects indicated the parallel approach traffic was deviating towards them, the traffic display was blanked and six possible maneuvers were shown to the subject: Turns in either direction; climbing turns in either direction, a straight-ahead climb, and a continued descent. The subjects were asked to select the maneuver they considered best for maintaining inter-aircraft separation.
- *Numerical Simulation of All Avoidance Maneuvers* The simulator then calculated the future trajectory of the intruder and of the subject aircraft for all six possible avoidance maneuvers. Performance metrics of each avoidance maneuver, such as the resulting miss distance, were calculated and stored, simplifying the data reduction process. These numerical simulations were transparent to the subjects and did not provide any feedback to them of their performance.

The simulator runs with each subject lasted one hour, including briefing, practice runs, all experiment runs, and a debriefing. The briefing explained the displays, controls and procedures involved in the experiment. Subjects were allowed as many practice runs as they requested, and additional practice runs were given before experiment runs with each new display. After the experiment runs, subjective comments were solicited from the subjects about the displays and their alerting strategies.

4.2.2 Independent Variables

4.2.2.1 Displays

Five displays were tested. All were based on a moving map display like the baseline TCAS display in the previous experiment, with a top-down view, track-up orientation, iconic presentation of the other aircraft's positions and a text presentation of the other aircraft's altitudes. All features of the traffic display were updated once per second, an update rate feasible with current technology.

- *Baseline Display*: Emulated the current TCAS display, as shown in Figure 4.1a.
- *Approach Fan Display*: Added the reference indication of the parallel approach path, emulating the approach fan shown on the enhanced EHSI display tested in Experiment #1, as shown in Figure 4.1b.
- *Heading Display*: Added an indication of the other aircraft's heading, shown in Figure 4.1c.
- *Noisy Projection Display*: Added a graphic indication of the other aircraft's heading rate and projected position within the next 15 seconds, as shown in Figure 4.1d. The position projection was based on the noisy measurement of the other aircraft's bank, using bank variations typical of those found during localizer tracking,
- *Smooth Projection Display*: Added a graphic indication of heading rate and projected position within the next 15 seconds, as shown in Figure 4.1d. The position projection used exact knowledge of heading rate to give a smoother projection.

4.2.2.2 Workload

The subjects were told their primary task was to keep their wings level despite turbulence using a side-stick. To do this, an artificial horizon was available to them, drawn approximately three inches away from the edge of the traffic display. For an additional incentive, a prize was offered to the subject whose bank deviation from level was the smallest, averaged over all data runs.

The turbulence was set to two different levels: in the high workload case, it required almost all of the subjects' attention, while in the low workload case it required checking, on average, only once every two seconds. The subjects were not briefed on these qualities. The turbulence in bank was provided by a Markov model, with the frequency of state changes and the probability of state changes set for both the High and Low workload cases through preliminary subject runs.

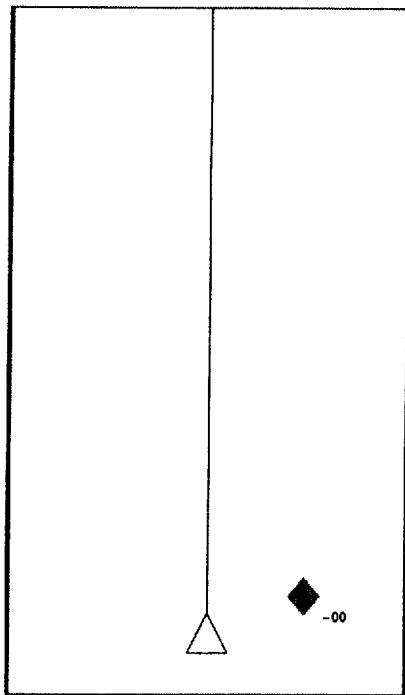


Figure 4.1a Baseline Display

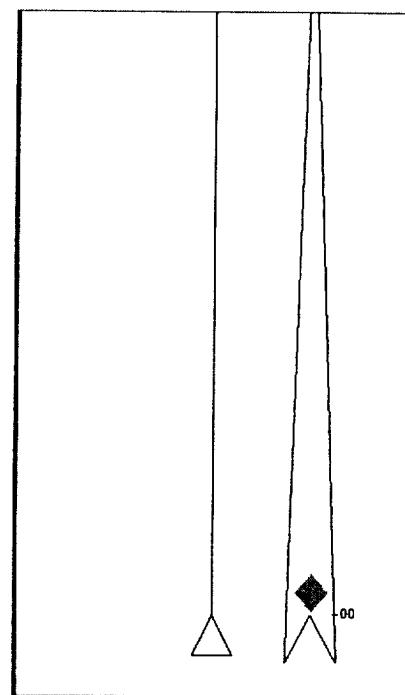


Figure 4.1b Approach Fan Display

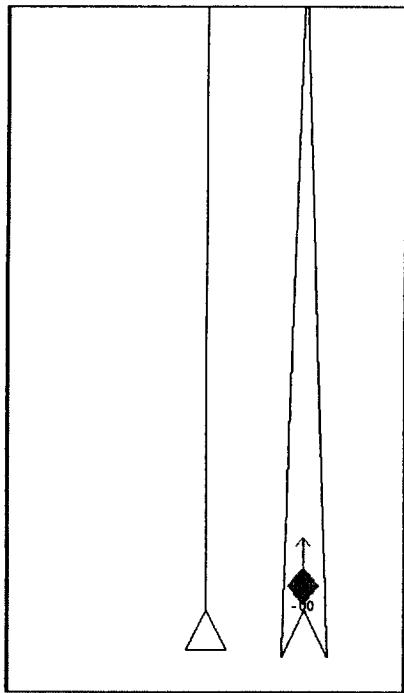


Figure 4.1c Heading Display

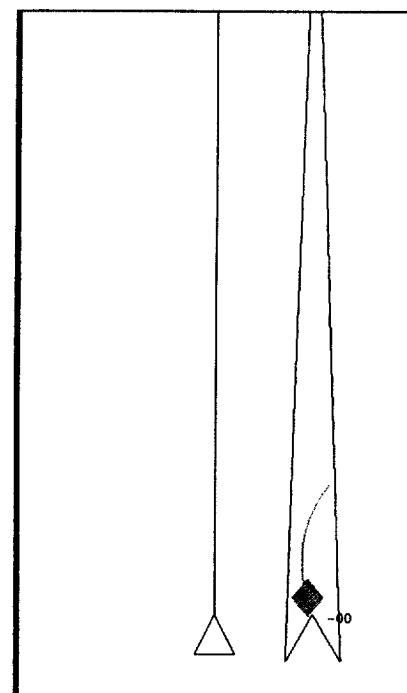


Figure 4.1d Projection Displays

4.2.2.3 Scenarios

Four scenarios were flown, in random order, within each test block. They were:

- *Missed Intercept*: The parallel traffic did not capture its parallel approach course but continued through its localizer intercept on a straight line collision course.
- *Hazardous Blunder*: The parallel traffic joined its own approach course, but at a random time during the approach turned towards the subject, establishing a collision trajectory.
- *Less Hazardous Blunder*: The parallel traffic joined its approach course, but at a random time during the approach blundered toward the subject, establishing a trajectory that passed at least 1000 feet away from the subject.
- *Safe Approach*: The parallel traffic flew its own approach normally, without deviating.

The intercept angle and turn rate were varied to create four possible conditions. The intercept angle of the intruding aircraft was picked to be high in one half of the runs with each display (45°) and to be low in the remaining one half of the runs (15°). Likewise, the turn rate of the intruder was set to a high value in one half of the runs ($4.5^\circ/\text{second}$) and low in the remaining half ($1.5^\circ/\text{second}$).

All cases represented parallel approaches 2000 feet apart. The intruder's speed was randomly selected at the beginning of each experiment run using a uniform distribution between 140 knots (the subject aircraft's speed) and 180 knots. The intruder was started at the correct altitude for glideslope intercept, and continued to follow the glideslope until it broke off its approach (or, in the Missed Intercept case, passes through its approach course) and deviated towards the subject. The vertical rate of the intruder during this blunder was also set in random order to each of three different values during each set of four runs. In one case, the intruder continued his approach descent; in another, the intruder used a $0.5g$ pull-up to level flight; and in the final case the intruder used a $0.5g$ pull-up to a climb of 2000 feet per minute.

4.2.2.4 Test Matrix

The test matrix for this experiment was three dimensional, varying displays, workload levels and traffic conflict scenarios. The test matrix is shown in Figure 4.2. Altogether, most subjects completed 40 experiment runs, allowing for within-subject comparisons; four subjects did not have runs with the smooth predictor display. The scenarios were flown in 10 blocks of four, where each block included all the runs for each particular display-workload combination.

		High Workload	Low Workload
		4 Scenarios	4 Scenarios
Baseline Display	+ Approach Fan for Other Runway	4 Scenarios	4 Scenarios
Baseline Display	+ Approach Fan for Other Runway	4 Scenarios	4 Scenarios
Baseline Display	+ Approach Fan for Other Runway	+ Other Aircraft Heading	4 Scenarios
Baseline Display	+ Approach Fan for Other Runway	+ Other Aircraft Heading	+ Noisy Projection
Baseline Display	+ Approach Fan for Other Runway	+ Other Aircraft Heading	+ Smooth Projection

Figure 4.2. Experiment #2 Test Matrix

4.2.3 Simulator Setup

The simulator used a Silicon Graphics Indigo 2 workstation for the displays and aircraft dynamics computations. A sidestick was connected for the flying task, and a mouse for the avoidance maneuver selection. The simulation was designed such that the subjects could easily control their progress, selecting further practice or commencement of the experiment runs.

The aircraft dynamics used simple point-mass calculations with performance constraints representative of air transport aircraft. The pitch steering and heading acquisition models used critically damped controllers, while the localizer acquisition controllers were slightly overdamped, modeling the actual wavering of the aircraft about the approach path.

4.2.4 Subjects

In total, nineteen subjects flew the experiment. Two were current airline flight crew, four were current Certified Flight Instructors (CFI) in general aviation aircraft (one with jet fighter experience), two held Private Pilot Licenses, and the remaining eleven were undergraduate or graduate students without piloting experience.

4.2.5 Measurements and Performance Metrics

Several measurements were taken, including:

- Type of scenario and variables defining the collision trajectory in each run
- Error at the workload task in each run
- Time and aircraft states when subject reacted in each run
- Avoidance maneuver selected by subject, and time required to select it in each run
- Subjective comments at the end of the experiment

Numerical simulations at the end of each run also projected the miss distance achieved with each of the six available avoidance maneuvers. While this does not provide an exact replication of the miss distance achieved by pilots manually controlling the aircraft, it does provide an approximate measurement of the subjects' decision making and of the timeliness of their reactions.

The statistical significance of differences between measurements made under different conditions were tested using unpaired t-tests.

4.3 Experiment Results

The results of this experiment will be discussed as follows. First, the overall performance of the subjects is discussed. Then the comparative effects of each of the elements in the test matrix (collision geometries, displays, and workload) are examined. Finally, the performance of subjects with different characteristics is discussed.

4.3.1 Overall Performance & Characteristics of Subject's Alerts

Several measures exist for examining the validity and speed of the subjects' determination that the intruder was deviating towards their own approach path. First, as shown in Table 4.1, the correctness of the subjects' reactions can be evaluated. Table 4.1 is divided into two parts, one

showing the correctness of the decisions for the scenarios where the traffic blundered off its approach path, the other showing the decisions for the “Safe Approach” scenarios where the traffic did not blunder but instead maintained its approach path. A Near Miss is defined as a miss distance less than 500 feet. (Carpenter & Kuchar, 1996)

In the cases where the intruder was scripted to deviate, the subject correctly spotted the anomaly in time to safely avoid the other aircraft almost 82% of the time. However, some deviations were not noticed by the subjects until the aircraft spacing was less than 500 feet (0.7% of the runs), and others were noticed so late that all of the six available avoidance maneuvers were projected to result in a Near Miss (16.4%). Subjects also reacted, in some runs, while the intruder was maintaining its proper approach course (1.0% of the runs). This classification of reactions only evaluates the timing of the reaction, and does not consider the possible effectiveness of an ensuing avoidance maneuver.

In the cases where the intruder was not scripted to deviate, the subjects correctly did not indicate a need for a reaction over 97% of the time. False alarms were given in 2.9% of the cases.

Table 4.1 Correctness of Subjects' Reactions

Type of Reaction	Description	% of Approaches
Scenarios Involving Intruder Deviation Towards Subject (n = 546)		
Correct Detection	A impending collision was spotted in time for an effective avoidance maneuver.	81.9%
Missed Detection	A Near Miss was spotted only after it has occurred	0.7%
Late Alert	A reaction was given before a Near Miss, but too late safely avoid the other aircraft	16.4%
Early Alert	The subject reacted while the intruder was still maintaining a correct approach course	1.0%
Scenarios Without an Intruder Deviation (n = 182)		
Correct Rejection	Subject correctly choose to not react.	97.1%
False Alarm	The subject reacted although the intruder maintained their correct approach course	2.9%

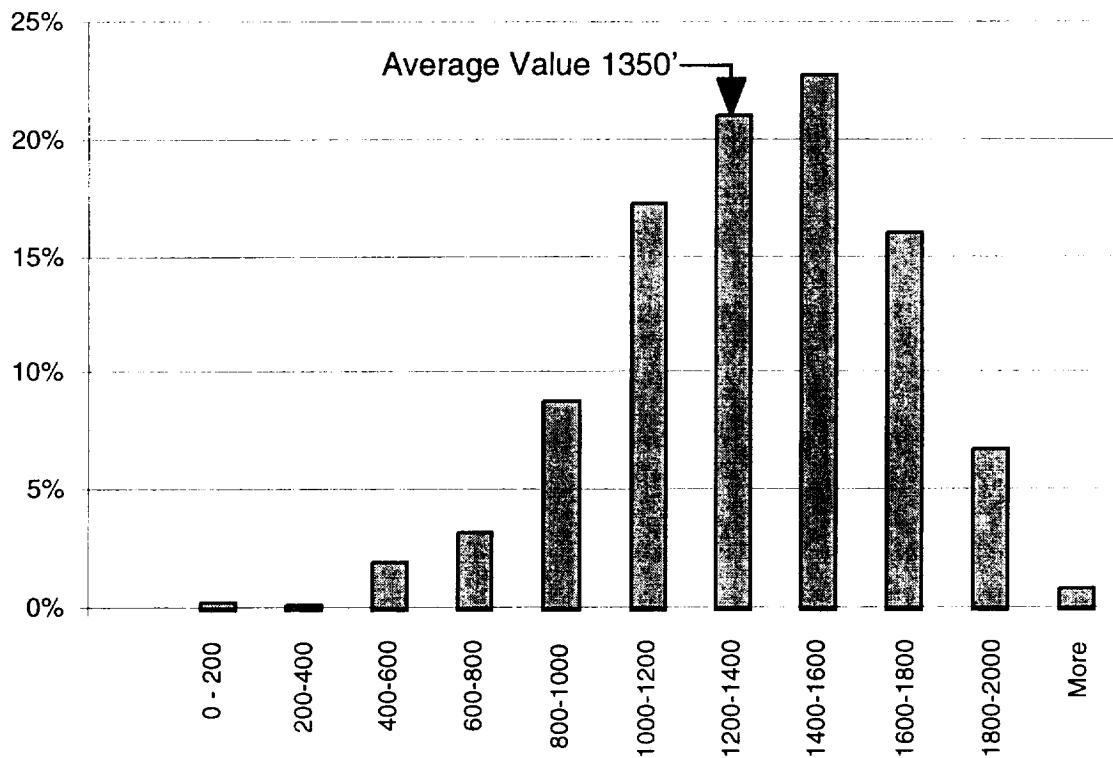


Figure 4.3 Histogram of Lateral Separation (Feet) When the Subjects Reacted (n = 546)

The histogram of the lateral separation between the aircraft at the time of the reaction is shown in Figure 4.3. A Chi-Squared goodness-of-fit test found its distribution approximates a normal distribution with a high probability ($p > 99\%$). The mean lateral separation at the time of the reaction is 1346 feet, with a standard deviation of 345 feet. For comparison, in the high convergence rate blunders, the aircraft lateral separation could decrease 200 feet between every one second update of information about the other aircraft. Therefore, the variance of this distribution is comparable to that expected from a standard deviation of 1.75 seconds in reaction time around an alerting criteria based purely on lateral separation.

The estimated time left until the point of closest approach was scattered, as shown in the histogram in Figure 4.4. The time remaining to point of closest approach ranged from -13.39 seconds (the subject reacted after the point of closest approach) to 34.32 seconds, with a mean of 14.37 seconds. The wide spread suggests the subjects' alerting criteria did not take into account convergence rate, differing from the alerting criteria used by TCAS II which uses convergence rate to estimate time remaining to collision and alerts once this time estimate is below a threshold value.

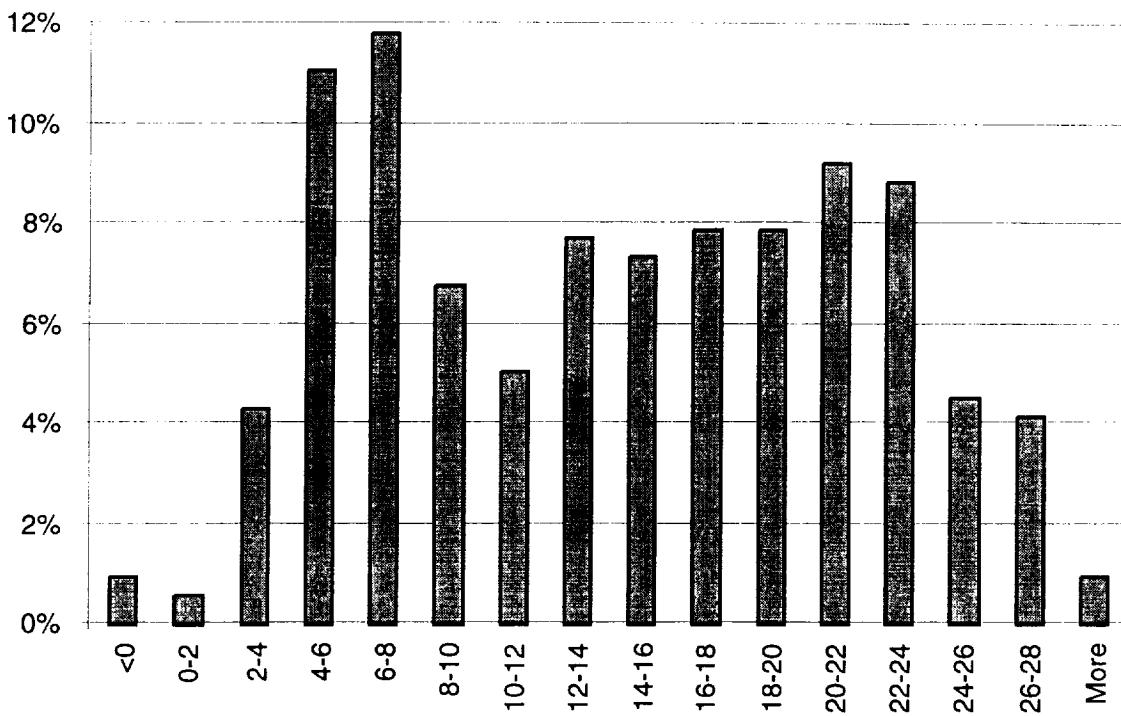


Figure 4.4 Histogram of Seconds Remaining to Point of Closest Approach When the Subjects Reacted (n = 546)

In addition to the timing of the subject's reactions, the performance of the subjects in selecting a safe direction of flight for an avoidance maneuver can be evaluated. The subjects were asked to select one of six possible avoidance maneuvers, and the performance of all six were projected.

Table 4.2 lists the frequency with which subjects selected each avoidance maneuver. The most popular maneuvers were Turn Away and Climb, and Turn Away (Maintaining Altitude), showing a strong preference for turning-away maneuvers. The remaining maneuvers were selected rarely.

Table 4.2 Frequency of Avoidance Maneuvers Selection (n = 642)

Maneuver	Frequency Maneuver Was Selected
Turn Away from Intruder (Maintain Altitude)	36%
Turn Away from Intruder and Climb	55%
Climb (No Turns)	3%
Turn Towards Intruder and Climb	2%
Turn Towards Intruder (Maintain Altitude)	2%
Continue the Approach	2%

The subjects' selected avoidance maneuvers were projected to cause Near Misses in 23% of the experiment runs. If the subjects had not reacted but instead had continued on their approach, Near Misses would have resulted in 43% of the approaches, and if they had chosen maneuvers randomly, Near Misses would have resulted in 38% of the approaches. Therefore, the subjects' selected avoidance maneuvers caused a significant improvement in collision avoidance, within the constraints of the limited choices in avoidance maneuvers.

4.3.2 Characteristics of the Subject's Reactions and Avoidance Maneuvers for Different Collision Geometries

Altogether, six different collision trajectories were tested. Each presented a different picture on the traffic display, and each caused a traffic hazard within a different amount of time. The convergence rate between the aircraft was controlled by setting the intruder intercept heading angle. In half of the experiment runs the intercept angle was set to be high (45°) and in the remaining half the intercept angle was set to be low (15°). Within the two Blunder scenarios, where the intruder's rate of turn from its approach path was also a factor, the heading rate was set to high and low values ($4.5^\circ/\text{second}$ and $1.5^\circ/\text{second}$ respectively).

Table 4.3 shows the length of time from the moment when the intruder left its nominal approach path to the time when the intruder crossed the subject's approach path. The times varied slightly because the intruder's speed for any approach was randomly varied between 140 knots and 180 knots. In the Missed Intercept scenarios, this measurement started when the intruder crossed the centerline of its approach path. In the Blunder scenarios, this measurement started when the intruder first started its turn towards the subject. Although this measurement is not necessarily the exact time to point of closest approach, it can be considered representative of the overall time-scale of the subjects' task.

Table 4.3 Potential Times to Collision With Different Scenario Variables

		Intercept Angle = 15°	Intercept Angle = 45°
Intruder Never Intercepts Own Approach (Missed Intercept Scenarios)		32.7 - 25.4 Seconds	12.0 - 9.3 Seconds
Intruder Intercepts Own Approach, Then Deviates (Blunder Scenarios)	Turn Rate = $1.5^\circ/\text{Sec}$	37.7 - 30.4 Seconds	26.2 - 23.5 Seconds
	Turn Rate = $4.5^\circ/\text{Sec}$	34.3 - 27.1 Seconds	16.7 - 14.0 Seconds

The correctness of the subjects' reactions are shown in Table 4.4. All of the Missed Detections and all but one of the Late Alerts occurred during the high-convergence Missed Intercept type scenarios. With the other collision geometries, the frequency of Correct Detections was very high. However, an alert was labeled a Correct Detection when it was given early enough that at least one of the six available maneuvers was projected to result in a safe miss distance; this categorization does not guarantee that the maneuver chosen by the subject was safe. The Correct Detection cases are also subdivided to show the frequency with which the maneuver selected by the subject resulted in safe and unsafe miss distance. Especially in the high convergence, high turn rate blunders, unsafe maneuvers were frequently chosen after a timely alert. This highlights the need for both correct detection of a traffic hazard, and then the selection of a safe avoidance maneuver.

The subjects tended to react in less time to the high convergence rate intrusions. However, these types of intrusions occur more quickly than the lower convergence rate cases. The difference in reaction time was not large enough for the reactions to be given with a common time remaining to collision. Instead, as shown in Figure 4.5, the time remaining to collision appears to have been largely a function of the collision geometry. Reactions were given for the high convergence rate cases with very little time remaining to collision. In contrast, reactions were generally given for lower convergence rate cases with a much greater time remaining to collision.

Table 4.4 Correctness of the Subjects' Reactions to the High and Low Convergence Intrusions

	Correct Detection		Missed Detection	Late Alert
	Maneuver Safe	Maneuver Unsafe		
Blunder With High Convergence Rate and High Turn Rate (n=90)	99%		0%	1%
	75%	24%		
Blunder With High Convergence Rate and Low Turn Rate (n=88)	100%		0%	0%
	94%	6%		
Blunder With Low Convergence Rate and High Turn Rate (n=88)	100%		0%	0%
	98%	2%		
Blunder With Low Convergence Rate and Low Turn Rate (n=89)	100%		0%	0%
	94%	6%		
Missed Intercept With High Convergence Rate (n=89)	4%		4%	91%
	2%	2%		
Missed Intercept With Low Convergence Rate (n=91)	100%		0%	0%
	95%	5%		

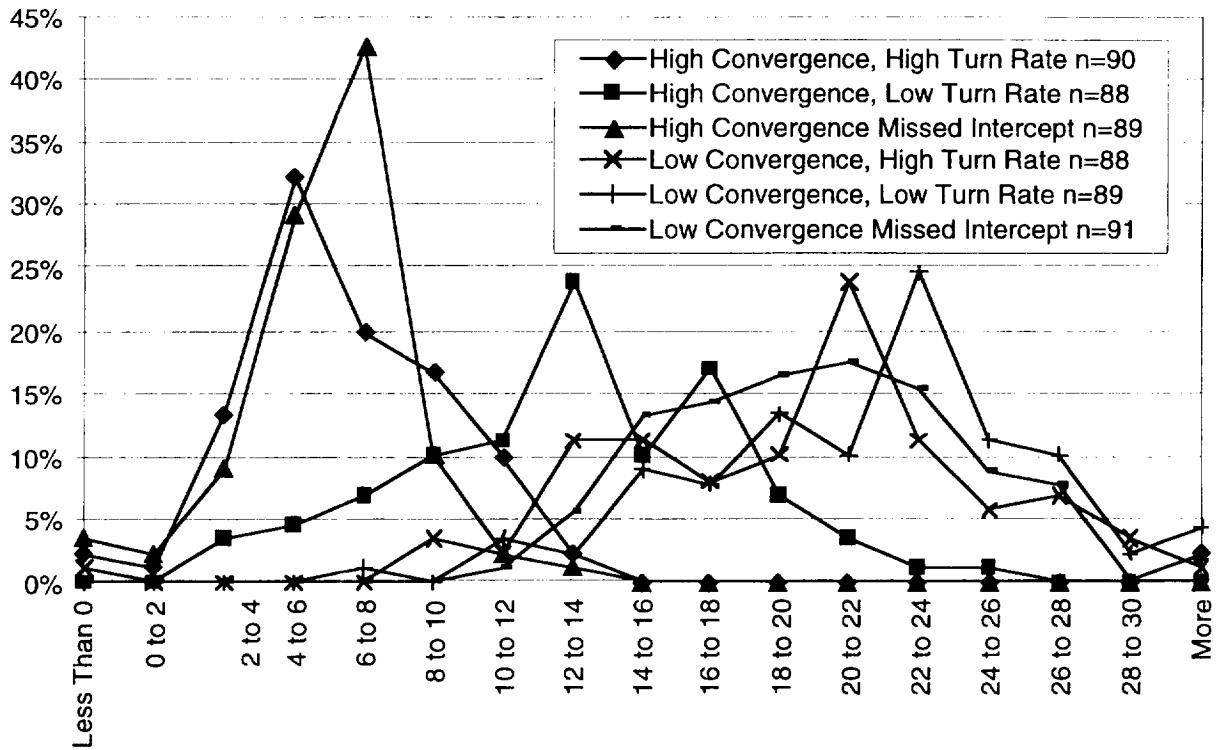


Figure 4.5 Time Remaining to Collision When the Subject Reacts, for Different Collision Geometries

Analysis of the subjects' reactions found that the same mean lateral separation existed between the aircraft at the subjects' reactions for the different collision geometries. The lateral separation at the subjects' alerts are shown in Figure 4.6. There is no statistically significant difference between them, except for the high convergence rate Missed Intercept scenario, which evolved very quickly and resulted in a lower lateral separation at alert than the other cases, even though the subjects' reaction times were the quickest to this type of collision trajectory.

The common lateral separation at subjects' reactions to different collision geometries suggests the pilots used lateral separation as a primary alerting criterion. The earlier reaction times with the high convergence intrusions appear to have been caused by the intruder reaching the critical lateral position earlier.

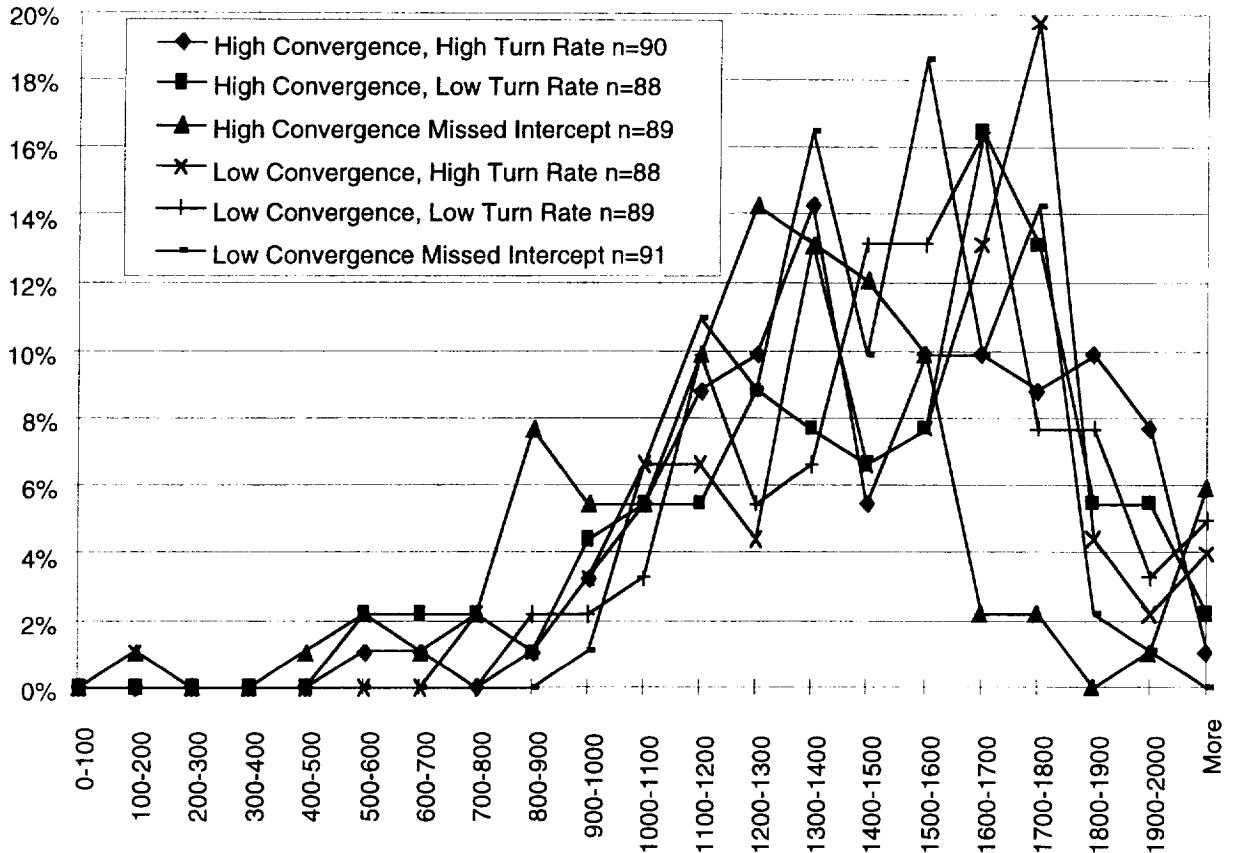


Figure 4.6. Aircraft Lateral Separation at the Subjects' Reactions, for Different Collision Geometries

4.3.3 Display Effects

In this experiment, five different displays were tested. Some display effects were found, and will be described using the same metrics as for the overall results. First, the correctness of the subjects' reactions is summarized in

Table 4.5. All of the False Alarms generated during the non-blunder scenarios occurred with the Baseline Display and with the Noisy Projection Display. All of the Early Alerts generated before a blunder started occurred with the Noisy Projection Display, a significant difference ($p < 0.05$). Runs with both the Baseline and Approach Fan displays each resulted in one Missed Detection, while the runs with the Noisy Projection display resulted in two Missed Detections, a difference that is not statistically significant.

Table 4.5 Correctness of Subject Reactions With the Different Displays

Display Type	Intruder Deviates Toward Subject				No Intruder Deviation	
	Correct Detection	Missed Detection	Late Alert	Early Alert	Correct Rejection	False Alarm
Baseline n=113	82.3%	0.9%	16.8%	0%	94.4%	5.6%
Approach Fan n=110	83.3%	0.9%	15.8%	0%	100%	0%
Aircraft Heading n=114	83.3%	0%	16.7%	0%	100%	0%
Noisy Projection n=114	77.2%	1.8%	15.8%	4.4%	91.7%	8.3%
Smooth Projection n=89	84.3%	0%	15.7%	0%	100%	0%

Although the newer displays were purposefully designed to give indications of relative convergence rate and trend, few differences can be found in the characteristics of the subjects' reactions with each of the different displays. Very little correlation between the inter-aircraft range and convergence rate can be found at the time of the subjects' reactions, indicating subjects did not tend to alert when a specific time to collision remained, even when shown indications of the intruder's trend.

Instead, the subjects' reactions appear consistent with a strategy using a low lateral separation from the intruder aircraft as a criterion for generating alerts. No statistically significant differences exist between the lateral separation at the subjects' reactions with the displays showing intruder heading and trend. The only statistically significant difference can be found in the lateral separation with the Baseline Display compared to the other displays; the lateral separation at the subject's reactions was significantly smaller with the Baseline Display ($p < 0.01$). This effect can be seen in Figure 4.7, and may be caused by the one difference between the Baseline Display and other displays -- the presentation of a reference of the other approach path. Commensurate longer reaction times were found with the Baseline Display.

Few significant differences can be found in the avoidance maneuvers selected by the subjects. Each maneuver appeared to be selected with the same frequency, regardless of display.

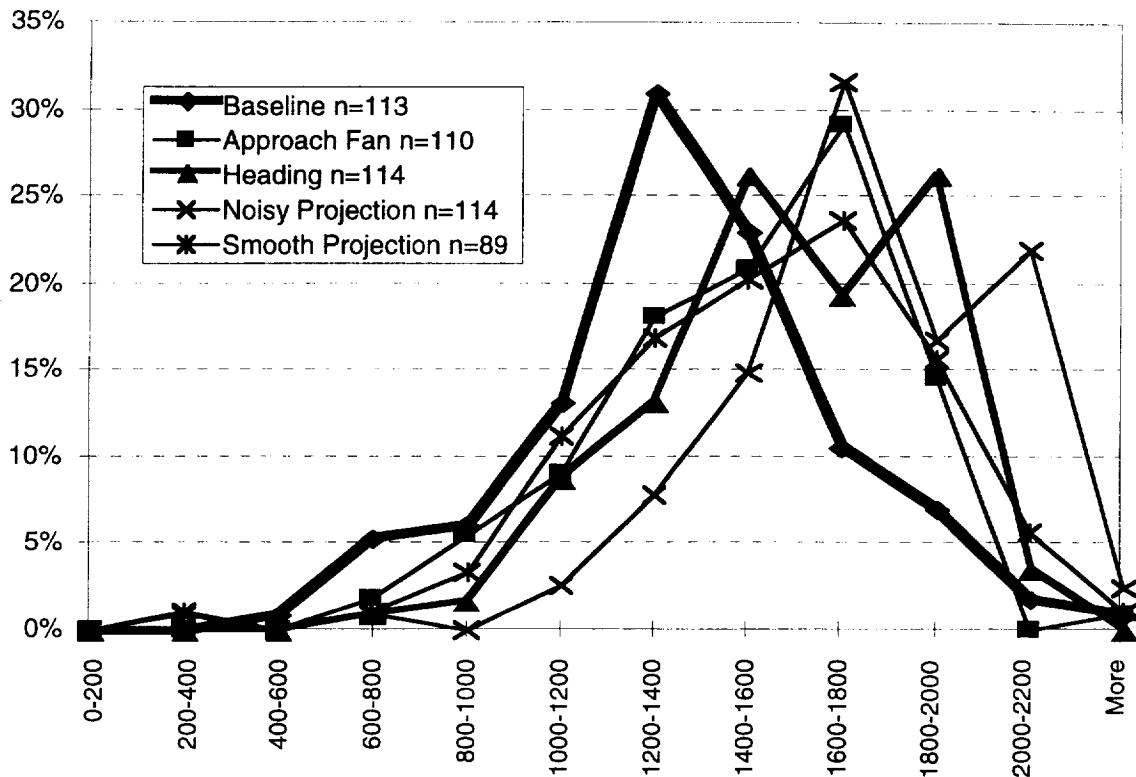


Figure 4.7 Aircraft Lateral Separation at Subject Reaction Times With Each Display

4.3.4 Workload Effects

Concurrent with monitoring a traffic display for possible traffic incursions, subjects were responsible for a wings-leveling task, using an artificial horizon and a side-stick. The difficulty of this task was manipulated to create a high or low workload for the subject to attend to away from the traffic display.

Most of the performance measures for this experiment were nearly identical when comparing the data from runs with high workload against runs with low workload. However, the Mean Squared Error (MSE) in the wings-leveling task itself was very different, as shown in Figure 4.8, especially for the period of time after the blunder had started. This illustrates the comparative difficulty of the different workloads.

The large difference between overall MSE bank and MSE bank once the blunder had started shows that, in high workload conditions, the subjects decided to drop the wings-leveling task in order to adequately assess the traffic situation. This differs from the subjects' briefings, in which

they were asked to consider the wings-leveling task to be primary. This increase may indicate the alerting decisions required a high level of attention.

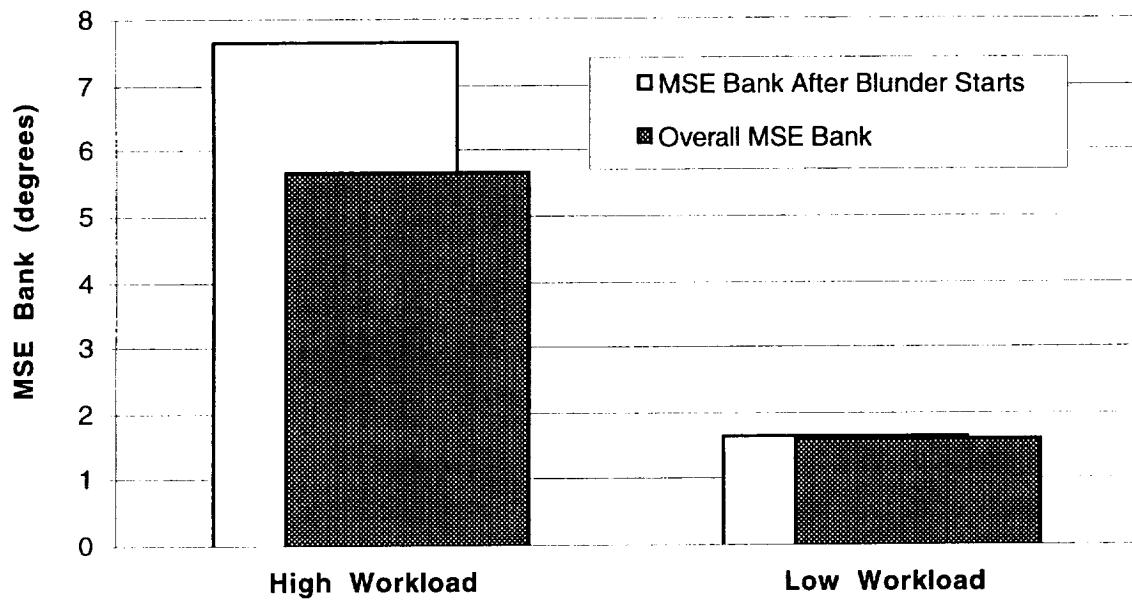


Figure 4.8 Mean Squared Error in the Sidetask for Different Subject Workloads

4.3.5 Variance With Subject Characteristics

Differences in performance between subjects of different characteristics were examined. Several measures were compared, including whether or not the subject was a pilot, and whether or not the subject described themselves as a “Video Game Junkie”.

With these comparisons, no statistically significant differences could be found in any of the standard performance metrics. The similar performance of pilots and non-pilots may suggest that this particular type of conflict detection is not something pilots are currently trained for, or accustomed to.

4.4 Summary of Results and Discussion

This experiment asked subjects to identify potential collisions without any form of automatic assistance other than displays of the intruder aircraft’s current state. The primary objective of this

experiment was to measure the characteristics of the subjects reactions. Secondary objectives examined the effect of displaying a reference of the parallel approach path and intruder trend information on the traffic display, and examined the effect of workload.

- Subjects' reactions appeared consistent with an alerting strategy using the lateral separation between themselves and the intruder as a primary criterion. With the display of a marker of the parallel approach path, subjects reacted with a slightly higher mean lateral separation. The display of heading and trend information did not appear to encourage different, more effective alert logic schemes.
- Subjects picked 'Turn Away' and 'Turn Away and Climb' avoidance maneuvers 91% of the time. These maneuvers did not always generate a safe aircraft separation. The type of maneuvers selected did not appear to be affected by the different traffic displays.
- Subjects tended to drop the workload-inducing sidetask once the intruder started deviating, indicating the task of monitoring traffic and deciding when to react may require a significant proportion of the pilots' attention at critical times.
- The similar performance of pilots and non-pilots as subjects in this experiment suggests that recognition of a potential traffic conflict during closely spaced parallel approaches is not a task pilots are accustomed to. Training may improve pilots' performance.
- In conclusion, the largest determinant of overall performance was the intruder's trajectory. Although subjects were effective at dealing with low convergence rate intrusions, the subjects' responses to high convergence rate intrusions generated the same performance as no reaction at all.

These results provide insight into the type of strategies used by subjects at these types of tasks. The subjects' reactions were consistent with simple criteria, based upon comparison of the position of fixed graphical fiducial markers on the traffic display. The display of more information about the other aircraft's trend did not promote the use of more accurate alerting strategies; this may have been caused by a lack of awareness by the subjects of their own low performance, or it may have been an indication that the subjects were not capable, in the time-available, of performing the extra computations required by more sophisticated strategies. Therefore, these results seem to emphasize a tendency of the subject to need a simple comparison upon which to base an alert.

This experiment suggests that subjects may disagree with the alerts and avoidance maneuvers made by more efficient alerting systems, such as the TCAS II system used in Experiment #1. These differences may contribute to non-conformance with the alerting system at critical times. Chapter 5 will discuss an experiment testing the use of more explicit justification of the alerting system decisions to reduce mismatches and increase pilot trust in the automatic decisions.

5. Experiment #3: Design and Evaluation of an Alerting System to Alleviate Problems with Pilot Non-Conformance

5.1 *Experiment Objectives*

Experiments #1 and #2 highlighted considerations in alerting system design for the task of collision avoidance during closely spaced parallel approaches. An alerting system is needed, as shown by the significant decrease in the rate of loss of aircraft separation when one was available. However, the full benefit of the alerting system was not realized due to pilot non-conformance. Pilots did not conform to a TCAS-type alerting system in 40% of the approaches. These non-conformance cases resulted in a significantly higher rate of loss of aircraft separation. Experiment #2 found subjects appear to use different criteria than alerting systems for reacting to a possible collision and selecting an avoidance maneuver. These differences in alerting and decision-making criteria may have contributed to pilots' non-conformance.

These results illustrate general considerations for the development of alerting systems which pilots will use effectively. In addition to the development of the alerting system's logic, the alerting system must be designed to encourage pilot conformance. This requirement may be more difficult when the alerting system's logic differs from that preferred by the pilots.

This experiment examines methods of promoting conformance to alerting system commands through the explicit display of the criteria underlying automatically generated alerts. The following conditions were tested:

- *No Automatically Generated Alerts, Criteria Explicitly Displayed* These experimental conditions tested the effect of the explicit display of alerting criteria on the subjects' reactions without the assistance of automatically generated alerts.
- *Automatically Generated Alerts, No Criteria Explicitly Displayed* These experimental conditions tested the effects of automatically generated alerts, in the absence of any explicitly displayed alert criteria.
- *Automatically Generated Alerts, Alert Criteria Explicitly Displayed* These experimental conditions were manipulated to create 'consonance' (where the explicitly displayed alert criteria was the basis for the automatically generated alerts) and 'dissonance' (where the automatically generated alerts were based on different criteria than that explicitly displayed to the subjects). In doing so, the effects of consonance and dissonance between the automatically generated alerts and the subjects' displays were examined.

5.2 Comparison of Alert Criteria for Closely Spaced Parallel Approaches

This experiment used two different alert criteria. The first alert criteria triggers at an inter-aircraft lateral separation of 1350'. This threshold was the mean lateral separation at the subjects' reactions in Experiment #2. This criteria is similar to the Non-Transgression Zone (NTZ) criteria used by the Precision Runway Monitor (PRM), as described in Appendix E. This criteria will be called the 'NTZ Alert Criteria' for the remainder of this experiment's discussion.

The 'MIT Alert Criteria' has been developed specifically for closely spaced parallel approaches. As such, this criteria is representative of the complex, higher-performance alerting strategies with which pilot conformance is desired. It is based on the probability contours shown in Figure 5.1. These contours delineate positions for an intruder aircraft, relative to the own aircraft, which have a specified probability of causing a collision, even if the own aircraft executes an avoidance maneuver that climbs and turns away from the intruder. The contour used as an alert threshold by the MIT Alert Criteria corresponds to the probability contour for $p = 0.001$; the alert criteria triggers when an intruding aircraft enters this area. (Carpenter & Kuchar, 1996)

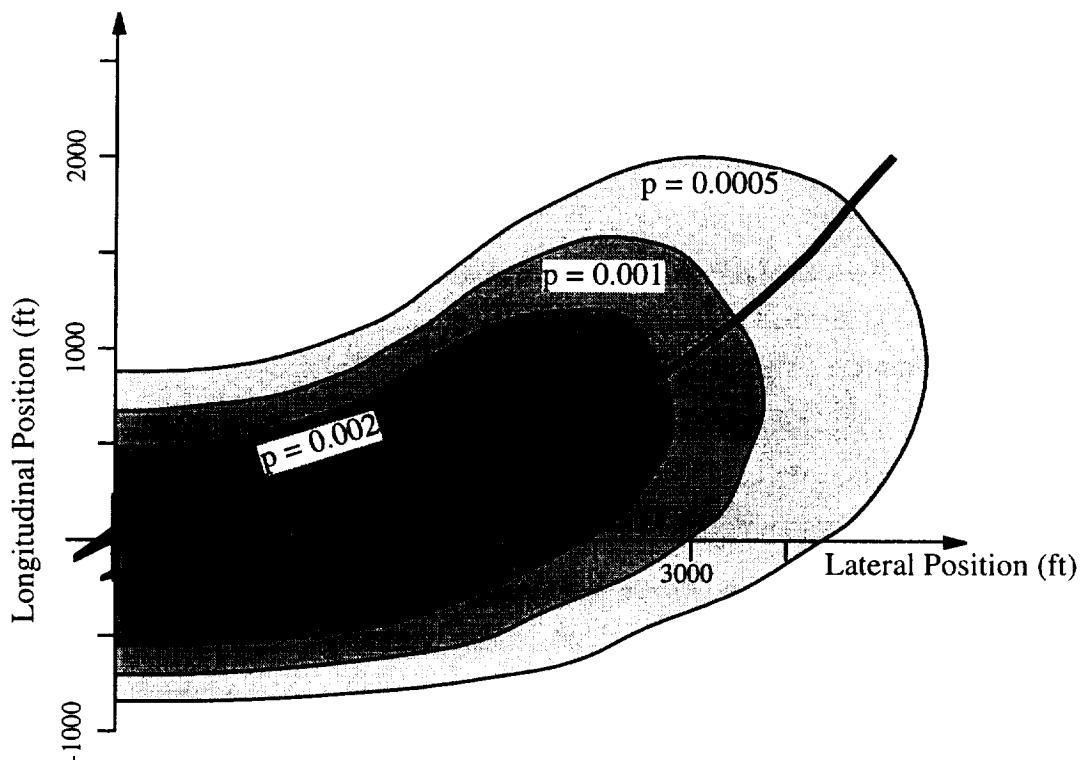


Figure 5.1 Probability of Collision Contours
(Figure From Carpenter & Kuchar, 1996)

The MIT criteria was developed assuming a climbing, turning avoidance maneuver will be executed once an alert has been triggered. This single, standard maneuver is the same as that preferred by subjects in Experiment #2. Therefore, it is expected pilots will be more likely to agree with the commanded avoidance maneuver (automatic decision). Although pilot conformance is likely to the automatic decisions commanded by this system, pilot conformance to automatic alerts -- as evidenced by initiating the avoidance maneuver at the time of the alert -- remains an issue. This experiment focuses on conformance to the automatic alerts.

As shown in Figure 5.2, the MIT criteria has two benefits over the NTZ criteria. First, intruding aircraft which will pass in front or behind the own aircraft trigger false alarms from the NTZ criteria, but not the MIT criteria. This is represented by Aircraft 1 in Figure 5.2, which will pass well in front of the own aircraft (shown at the origin).

Second, when the intruding aircraft poses a significant threat to the own aircraft, the MIT criteria triggers alerts in time for an effective avoidance maneuver. Alerts triggered by the NTZ alert criteria may be late, as illustrated by Aircraft 2 in Figure 5.2, which has established a high convergence rate and may cause a collision too quickly for an effective avoidance maneuver.

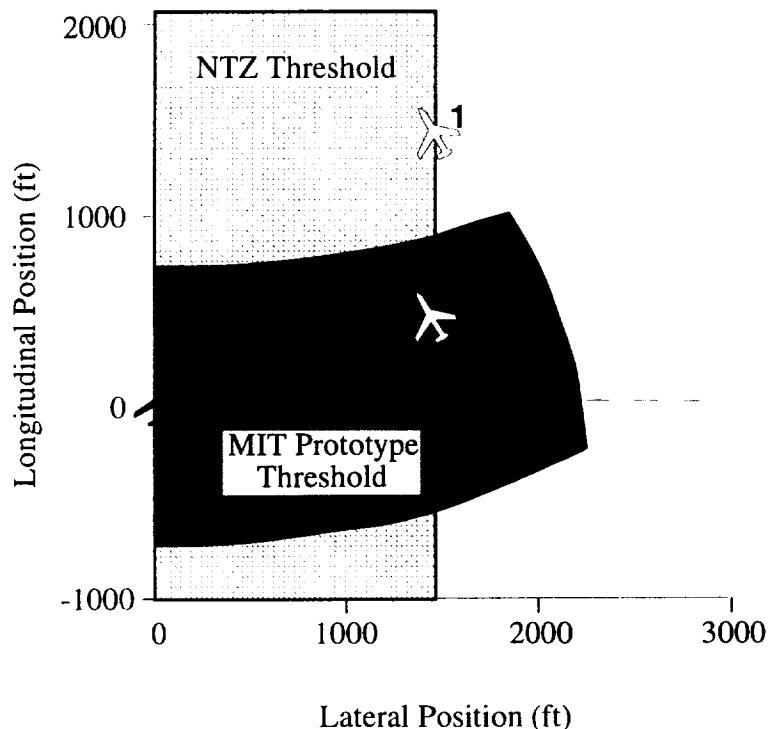
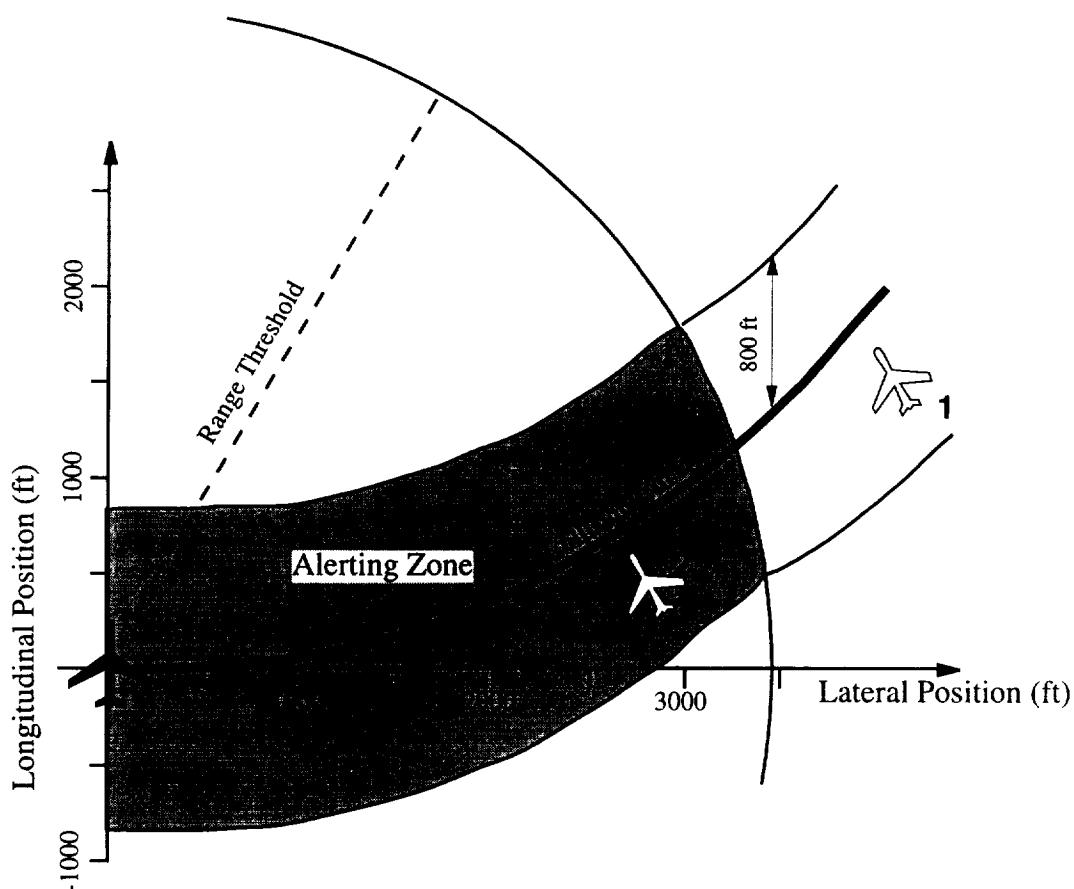


Figure 5.2 Comparison of the MIT and NTZ Alert Criteria
(Figure From Carpenter & Kuchar, 1996)

The MIT criteria can be approximated by a relatively simple geometry, as shown in Figure 5.3. The center of the area is delineated by a 'Collision Curve', which represents the relative positions from which an intruding aircraft, if it maintains its current bank and heading trend, will generate a collision. A margin of 800 feet is added fore and aft of the Collision Curve; the size of this margin was determined to be necessary during development to account for future changes in the intruder trajectory and variability in the sensor information. The range the alert boundary extends along the Collision Curve was also determined by Monte Carlo simulations during the design of the alerting logic. This range depends on the intruding aircraft's speed, heading and bank, and needs to be interpolated from a three dimensional lookup table.

Both the curvature of the Collision Curve, and the range that the alert boundary extends along it, vary with the other aircraft's heading, bank and speed. If drawn in real-time on a traffic display, the alert boundary changes with each new update of information about the other aircraft.



**Figure 5.3 Geometric Simplification of the MIT Alert Criteria
(Figure From Carpenter & Kuchar, 1996)**

5.3 Experiment Design

5.3.1 Experimental Procedure

Using a workstation based, part-task simulator, each subject completed 36 experiment runs. Each run consisted of three sequential parts:

- *The Flight* The subjects were told they were flying an approach. Their primary task was to keep their wings level despite turbulence through the use of a sidestick, referenced to an artificial horizon. The sidestick commands did not affect the path of the aircraft, so that consistent approach paths were followed. Their secondary task was to indicate when they thought the aircraft on a parallel approach was blundering towards them, as evidenced by the traffic display. Subjects pressed different buttons indicating whether they felt an avoidance maneuver was required by the traffic situation or not. The approach paths were separated by 2000 feet, the same runway spacing as in Experiment #2. In some cases, automatic alerts were given, indicated by both an aural alert and graphical indicators on both the artificial horizon and the traffic display. Subjects were told to use their best judgment in deciding when to react. Conformance to the automatic alerts was not mandated.
- *Certainty and Timeliness Ratings* Once the subjects indicated the parallel approach traffic was deviating, the traffic display was blanked and subjects were asked to indicate their certainty in their decision, using a mouse, on a graphical scale shown on the computer screen. If the subjects had been shown an automatic alert before their reaction, they were then also asked to rate the timeliness of the automatic alert on a similar rating scale.
- *Numerical Simulation of Avoidance Maneuvers* The simulator then projected the future trajectory of the intruder and of the subject aircraft throughout avoidance maneuvers triggered by the subject's reaction, by the NTZ alert criteria, and by the MIT alert criteria. Each avoidance maneuver commanded a climbing turn away from the intruder, and an increase in aircraft speed. Performance metrics of each avoidance maneuver, such as the resulting miss distance, were calculated and stored. These numerical simulations were transparent to the subject.

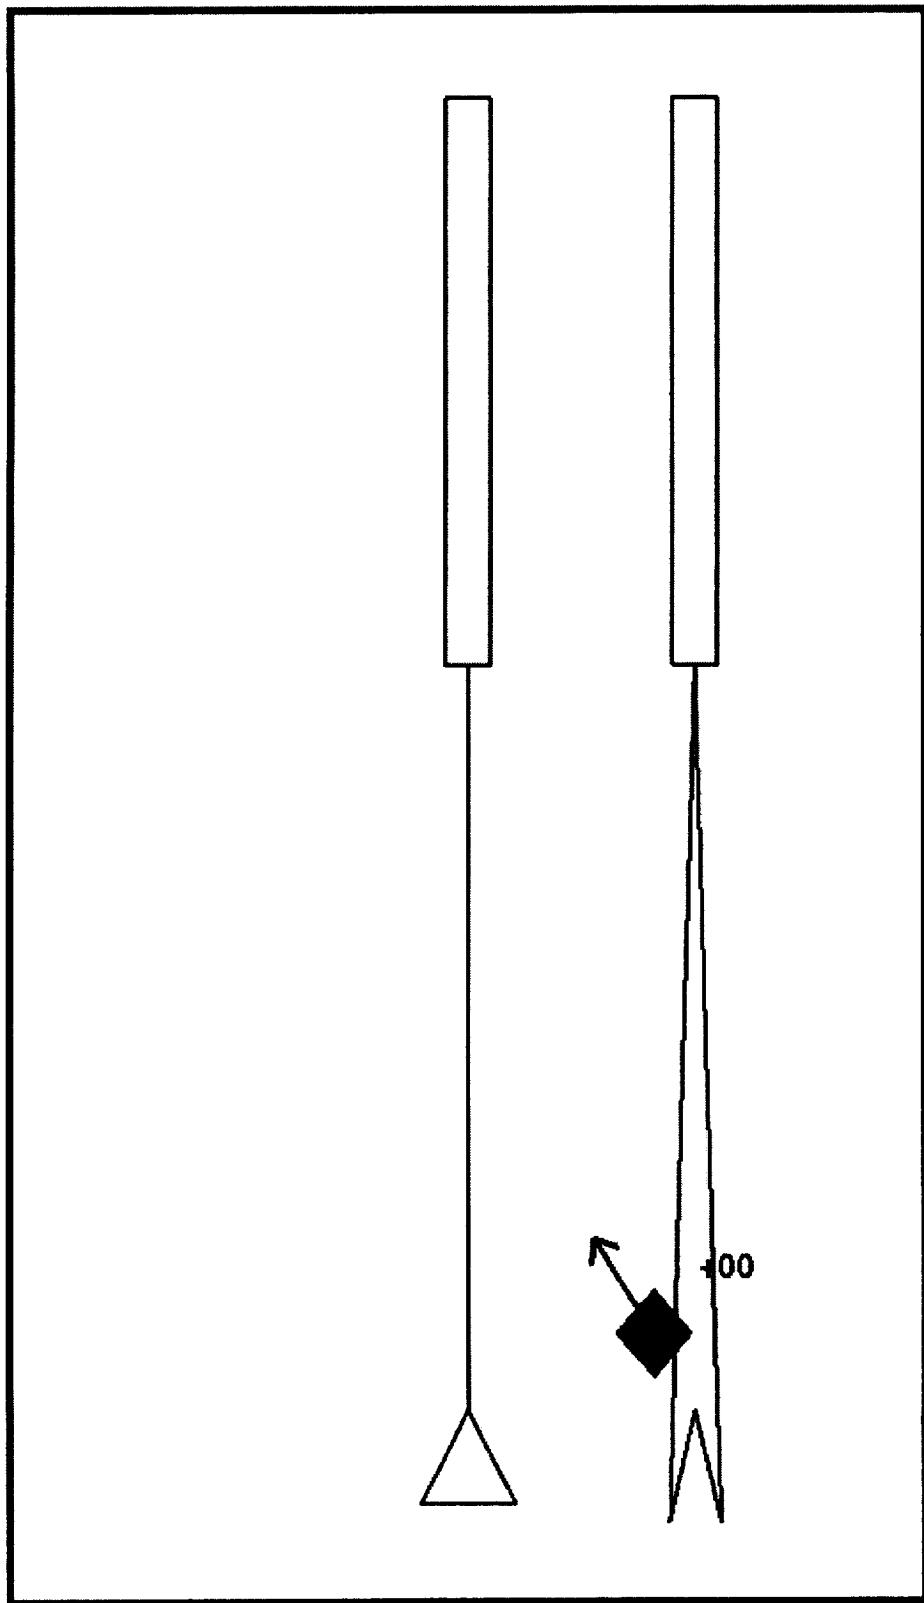
The simulator runs with each subject lasted one hour, including briefing, practice runs, all experiment runs, and a debriefing. The briefing explained the displays, controls and procedures involved in the experiment. Subjects were allowed as many practice runs as they requested. After the experiment runs, subjective comments were solicited about the displays and their alerting strategies.

5.3.2 Independent Variables

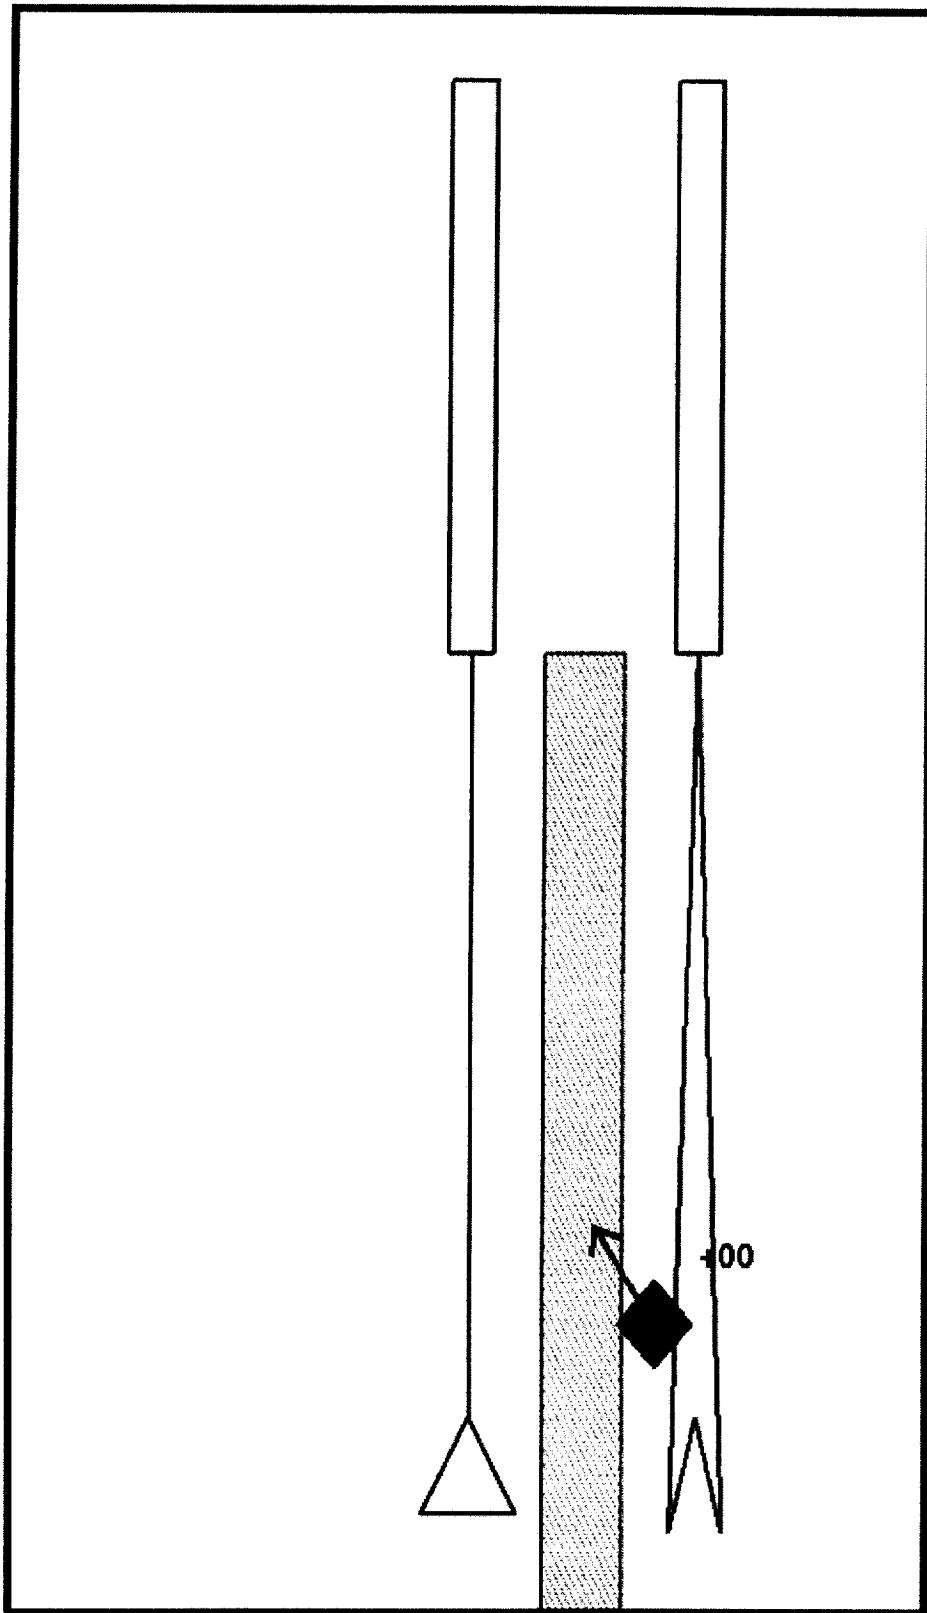
5.3.2.1 Displays

Three displays were tested. All were based on the moving map display used in Experiment #2, with a top-down view, track-up orientation, iconic presentation of the other aircraft's positions and a text presentation of the other aircraft's altitude. All features of the traffic display were updated once per second, an update rate feasible with current technology.

- *Baseline Display:* Emulated the current TCAS traffic display, with an additional indication of the other aircraft's heading, as shown in Figure 5.4. This was the subjects' preferred display in Experiment #2.
- *NTZ Criteria Display:* Added a graphic indication of a Non-Transgression Zone between the parallel approaches to the baseline display, as shown in Figure 5.5. The subjects' reactions in Experiment #2 were consistent with a NTZ type criteria. The dimensions of the NTZ criteria in this experiment were set to trigger an alert at the mean lateral separation used by the subjects in Experiment #2, 1350 feet.
- *MIT Criteria Display:* Added a graphic indication of the alert criteria used by the prototype MIT alerting logic to the baseline display, as shown in Figure 5.6. Because the shape of the alert criteria changes with each once per second update state of the other aircraft, the display of this alert criteria changed shape once per second. This effect is shown in three display snapshots in Figure 5.7, with the alert thresholds drawn relative to the fixed 'own aircraft' triangle at the bottom center of the screen.



**Figure 5.4 Baseline Traffic Display
(Grayscale, Black-White Inverted for Clarity)**



**Figure 5.5 NTZ Alert Criteria Shown on Traffic Display
(Grayscale, Black-White Inverted for Clarity)**

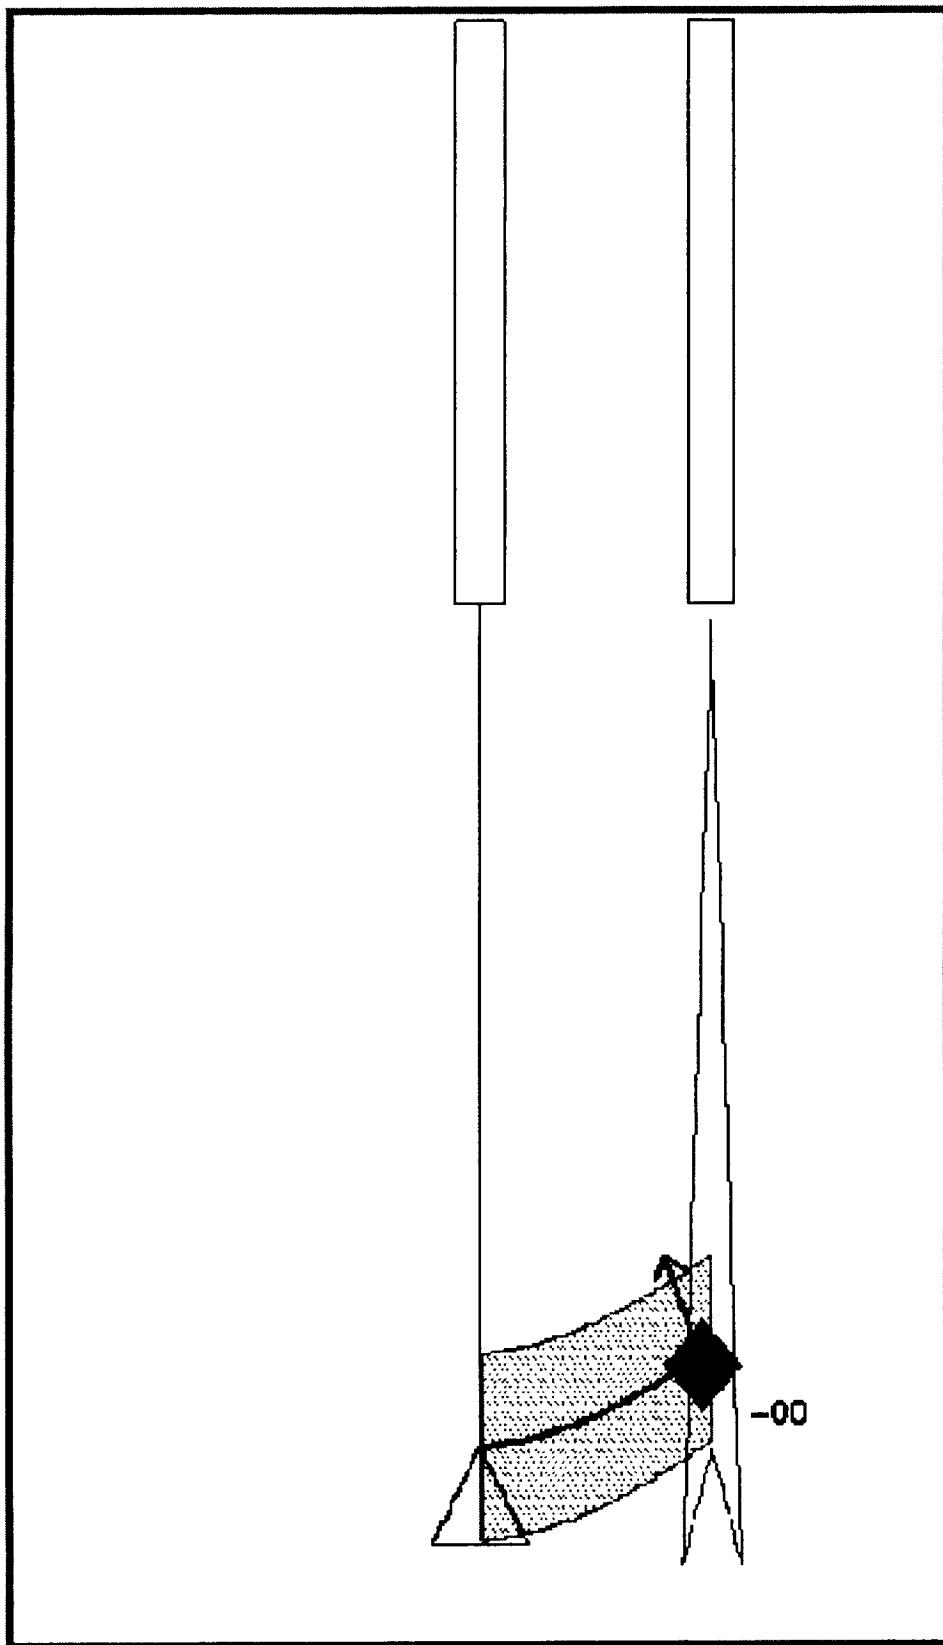
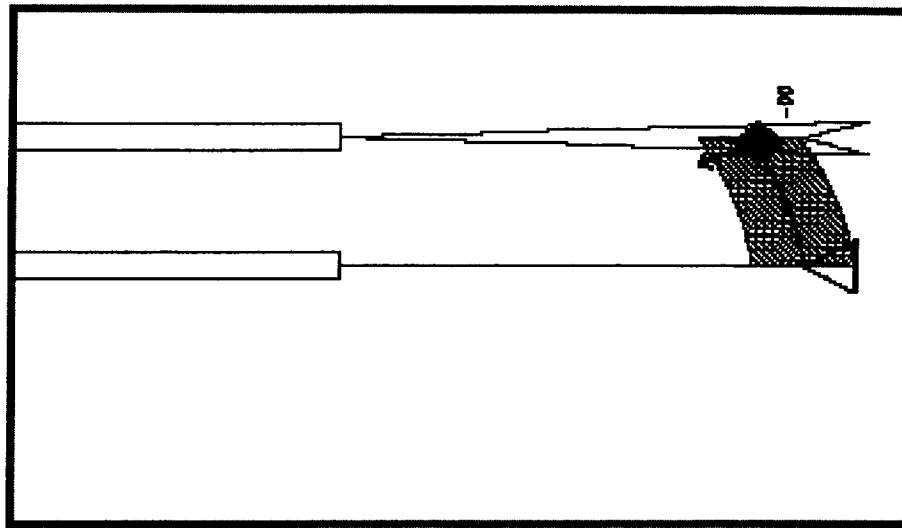
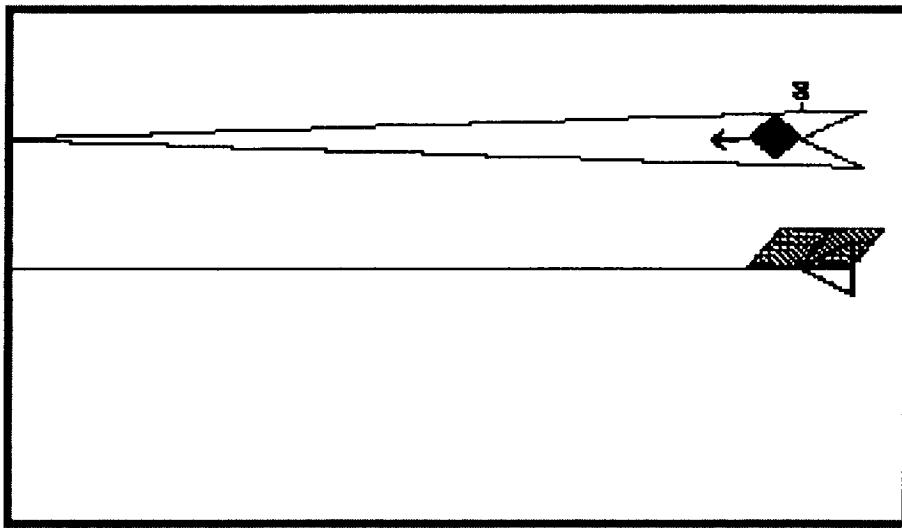


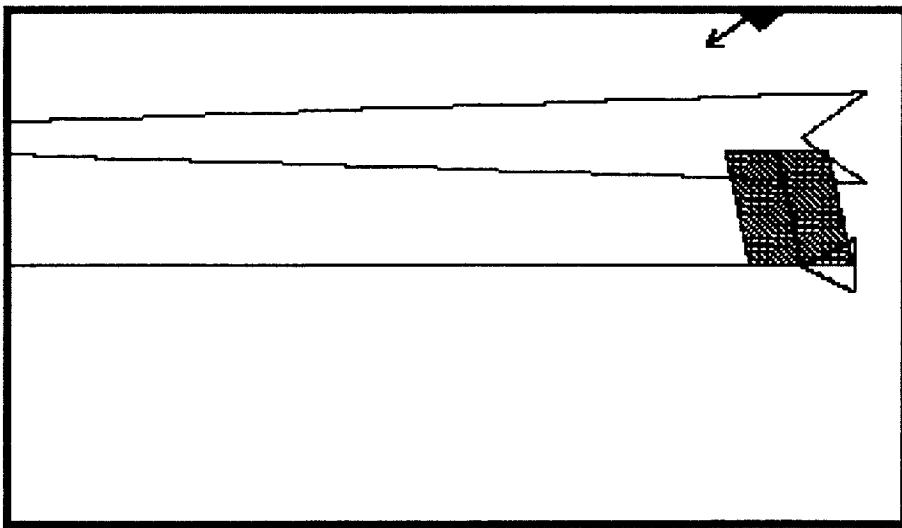
Figure 5.6 MIT Alert Criteria Shown on Traffic Display
(Grayscale, Black-White Inverted for Clarity)



Alert Criteria As Intruder Banks Towards Own Aircraft, At Start of Hazardous Blunder



Alert Criteria With Intruder Wavering Slightly Around Parallel Approach Course



Alert Criteria With Relative Intruder Heading of 35°, Level, During Intercept of Parallel Approach

Figure 5.7 Changes in Size and Shape of the MIT Alert Threshold With Intruder Aircraft Actions

5.3.2.2 Automatic Alerts

Three different automatic alerting conditions were used in the experiment:

- No automatic alerts were given to the subjects.
- Automatic alerts based on an NTZ criteria were given. This underlying criteria was the same as that shown explicitly on the NTZ Alert Criteria display.
- Automatic alerts based on the MIT prototype alerting logic were given. This underlying criteria was the same as that shown explicitly on the MIT Alert Criteria display.

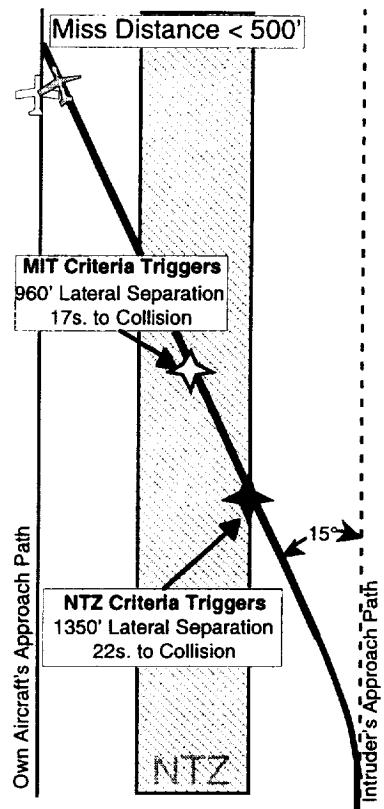
5.3.2.3 Scenarios

Four scenarios were flown, in random order, within each test block. They were:

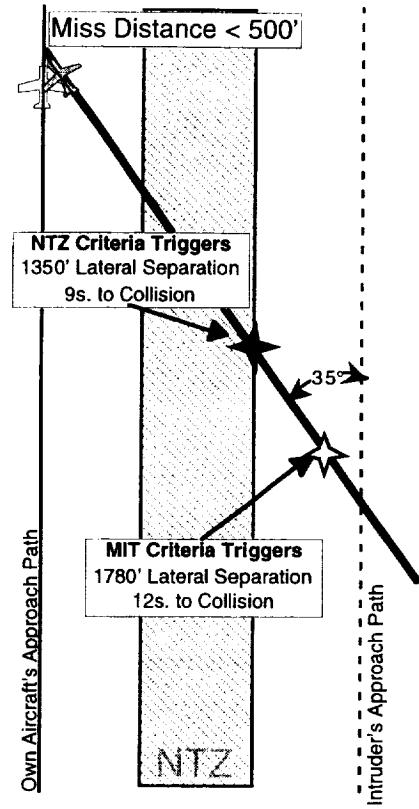
- *Low Convergence, Hazardous Blunder* The parallel traffic joined its own approach course, but at a random time during the approach turned towards the subject with a 15° bank to a 15° intercept heading, establishing a collision trajectory.
- *High Convergence, Hazardous Blunder* The parallel traffic did not capture its parallel approach course but continued through its localizer intercept on a straight line collision course, with an intercept heading of 35°.
- *Low Convergence, Less-Hazardous Blunder* The parallel traffic joined its own approach course. At a random time during the approach, it then turned towards the subject with a 20° bank to a 20° intercept heading, establishing a trajectory which passed more than 1000' feet away from the own aircraft.
- *High Convergence, Less-Hazardous Blunder* The parallel traffic joined its own approach course. At a random time during the approach, it then turned towards the subject with a 35° bank to a 35° intercept heading, establishing a trajectory which passed about 500 to 600 feet away from the own aircraft.

These scenarios are illustrated in Figure 5.8. The hazardous-blunder scenarios were designed to trigger the two different alert criteria at different times. In the low-convergence, hazardous-blunder scenario, the NTZ criteria triggered before the MIT criteria; in the high-convergence, hazardous-blunder scenario, the MIT criteria triggered first.

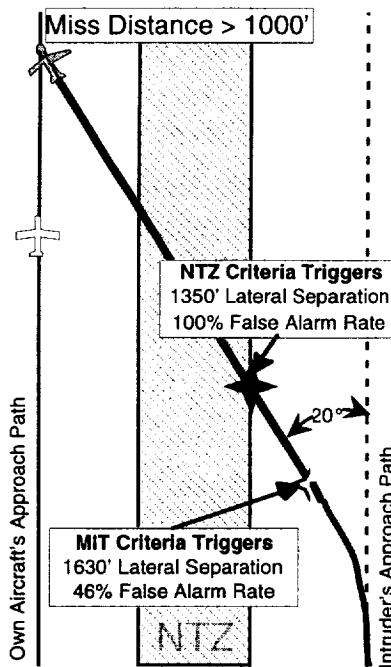
The NTZ criteria always triggered false alarms in the less-hazardous scenarios. The MIT criteria triggered false alarms about 46% of the time in the low-convergence, less-hazardous-blunder scenario, and almost 100% of the time in low-convergence, less-hazardous-blunder scenario, which resulted in miss-distances just over 500'.



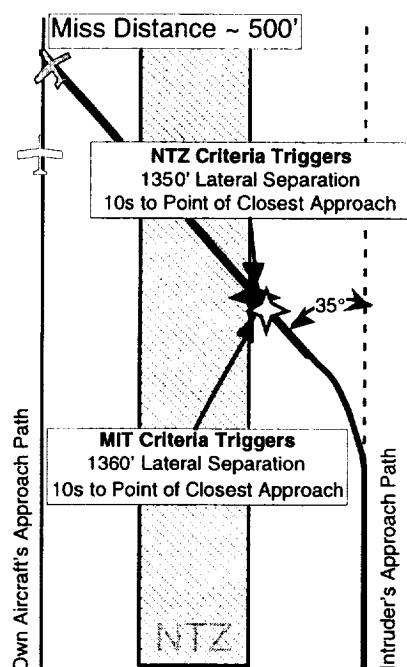
Low Convergence, Hazardous Blunder



High Convergence, Hazardous Blunder



Low Convergence, Less-Hazardous Blunder



High Convergence, Less-Hazardous Blunder

Figure 5.8 Illustration of the Four Different Scenarios

5.3.2.4 Test Matrix

The test matrix for this experiment was three dimensional, varying displays, alerts and traffic conflict scenarios. The test matrix is shown in Figure 5.9. Altogether, subjects completed 36 experiment runs, allowing for within-subject comparisons. The scenarios were flown in 9 blocks of four, where each block included all the runs for each particular display-workload combination.

	Baseline Display	NTZ Alert Criteria Shown	MIT Alert Criteria Shown
No Automatic Alerts Given	4 Scenarios	4 Scenarios	4 Scenarios
Automatic Alerts Based on NTZ Criteria	4 Scenarios	4 Scenarios	4 Scenarios
Automatic Alerts Based on MIT Criteria	4 Scenarios	4 Scenarios	4 Scenarios

Figure 5.9. Experiment #3 Test Matrix

The experiment runs tested a total of nine different alert-display combinations, as shown in Figure 5.10. These combinations created several different conditions of interest:

- When no automatic alerts were given, the effects of displaying alert criteria on subject decision-making were isolated;
- When no alert criteria was displayed, the reactions of the subjects to automatic alerts were isolated; and
- When both automatic alerts were given and alert criteria was displayed, the effects of the display being consonant with the alerts (the highlighted diagonal cases) and of the display being dissonant with the alerts (the off-diagonal cases) were examined.

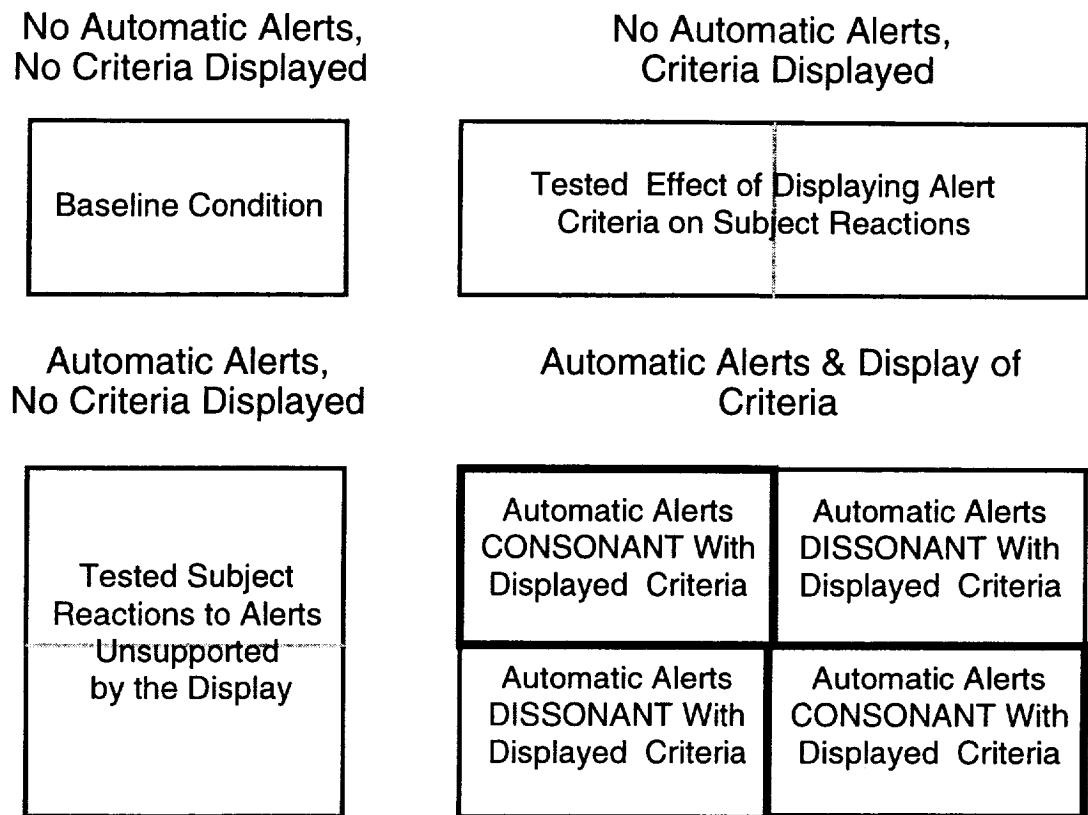


Figure 5.10 Expanded Test Matrix Showing Nine Alert-Display Combinations

5.3.3 Simulator Setup

The simulator used a Silicon Graphics Indigo 2 workstation for the displays and aircraft dynamics computations. A sidestick was connected for the flying task, and a mouse for the avoidance maneuver selection. The simulation was designed such that the subject could easily control their progress, selecting further practice or commencement of the experiment runs.

The aircraft dynamics used simple point-mass calculations with performance constraints representative of air transport aircraft. The pitch steering and heading acquisition models used critically damped controllers, while the localizer acquisition controllers were slightly overdamped, modeling the wavering of an actual aircraft about the approach path.

5.3.4 Subjects

In total, twelve subjects flew the experiment. Three held Certified Flight Instructor (CFI) ratings; six had some general aviation flight experience, and the remaining three subjects were graduate students without piloting experience. No subjects were, or had been, airline flight crew.

5.3.5 Measurements and Data Analysis

Several measurements were taken, including:

- Type of scenario and variables defining collision trajectory in each run.
- Time and aircraft states when subject reacted and the maneuver/ no maneuver decision in each run.
- Subjective ratings of certainty in their decision and, when appropriate, ratings of the timeliness of the automatic alerts in each run.
- Subjective comments at the completion of the experiment.

Numerical simulations at the end of each run also evaluated the miss distance projected from avoidance maneuvers triggered by the subjects' reactions and by each alert criteria. While this numerical simulation does not provide an exact replication of the miss distance achieved by pilots manually controlling the aircraft, it does provide a first order estimate of the timeliness of their reactions.

Based on the results of the numerical simulation, the subjects' reactions were categorized by whether the subjects decision was correct and timely enough to avoid a collision. Subjects' reactions were classified as Late Alerts if they identified a developing collision hazard too late for the projected avoidance maneuver to maintain the required 500' aircraft separation. Subjects' reactions were classified as False Alarms if they indicated a need for an avoidance maneuver when continuing the approach would have resulted in the required 500' aircraft separation.

Statistical tests were used to ascertain the significance of measures taken under different conditions. For continuous measures, paired t-tests were used. For discrete measures, paired sign tests were used.

5.4 Experiment Results

This section details the results of this experiment. First, the effects of the four different scenarios will be compared. Then, the effects of the displays and of the effects of automatic alerts will be examined. Finally, the combined effects of displaying alert criteria on subjects' acceptance of automatic alerts will be compared.

5.4.1 Effects of Scenarios

Four different scenarios were tested. As shown previously in Figure 5.8, these scenarios were designed to trigger the two different alerting criteria at different times. These differences in timing will be used throughout the analysis of the experimental results as a means of comparing the subject's decisions to those based on the two different alert criteria.

Examination of the subjective certainty ratings given by the subjects after each run found the subjects were significantly less certain about their decisions in the high-convergence hazardous blunder scenario, and significantly more certain in the low-convergence, less-hazardous blunder. These ratings may reflect the time the subjects had to consider their decisions in each of these two cases. Subjects quickly realized no penalty existed for delaying their decision when a collision was not projected to occur, and therefore often did not indicate the decision to continue the approach until the intruder was past them in the less-hazardous scenarios, possibly allowing for a higher certainty rating in these cases.

5.4.2 Display Effects

Three different displays were tested: a baseline display, a display explicitly showing the NTZ criteria, and a display explicitly showing the MIT criteria. This section discusses the effects of the displays on the type and timing of the subjects' reactions, aggregate over all alerting conditions.

Subjects' reactions tended to shift towards the criteria shown on the displays. In the hazardous-blunder scenarios, the time and lateral separation at subjects' reactions were closer to when the alert criteria shown on the display would have triggered. This effect is shown in the timelines for the two hazardous scenarios in Figure 5.10 and 5.11.

The overall differences between reactions with the baseline display and with the NTZ criteria display in both time and lateral separation were statistically significant to the $p < 0.01$ level. The difference in lateral separation between the baseline display and the MIT criteria display was also

significant to the $p < 0.01$ level; however, the difference in alert time between these two displays only tested significant to the $p < 0.10$ level.

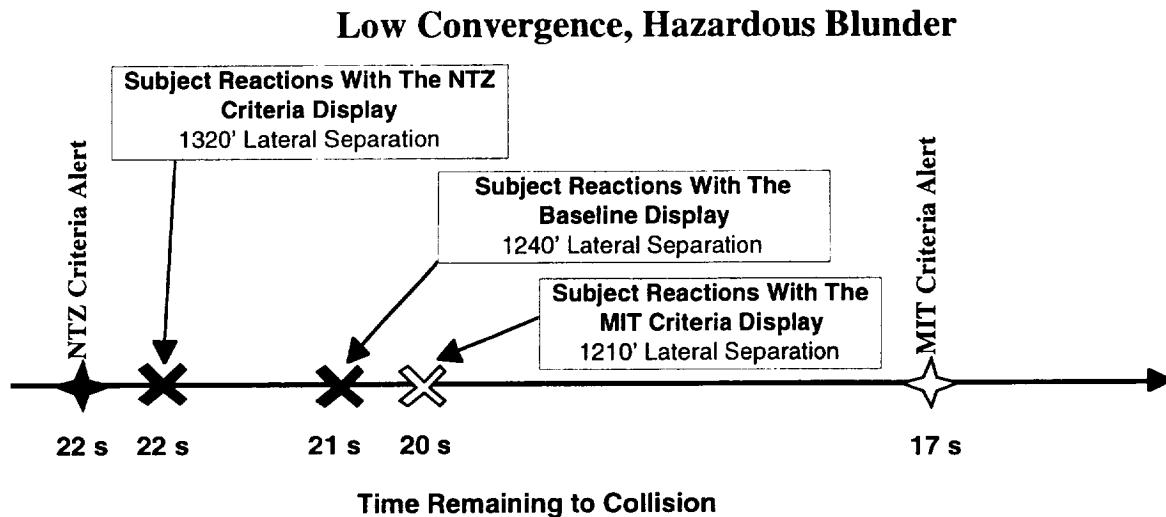


Figure 5.11 Timing of the Subjects Reactions in the Low Convergence, Hazardous Blunder Scenario, By Display

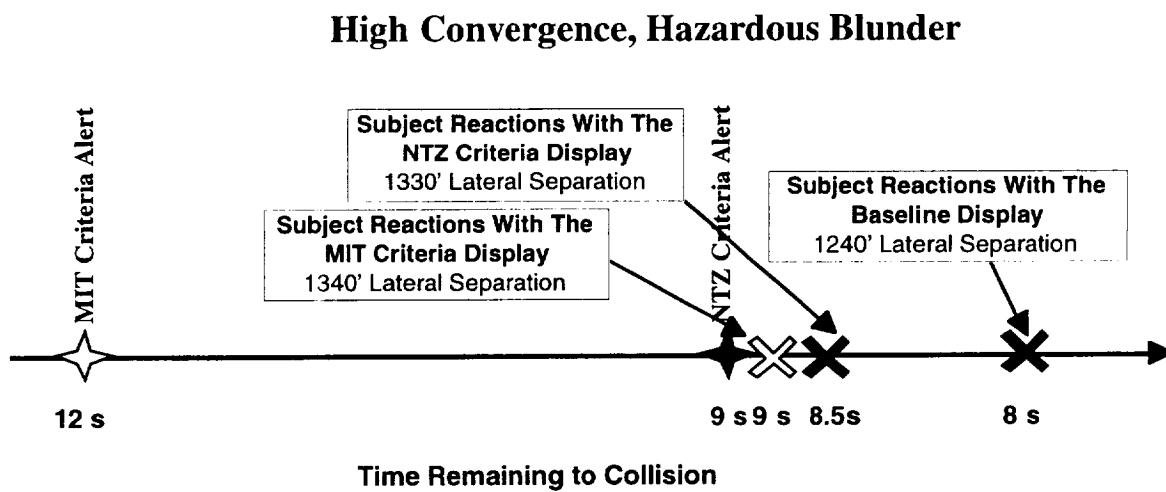
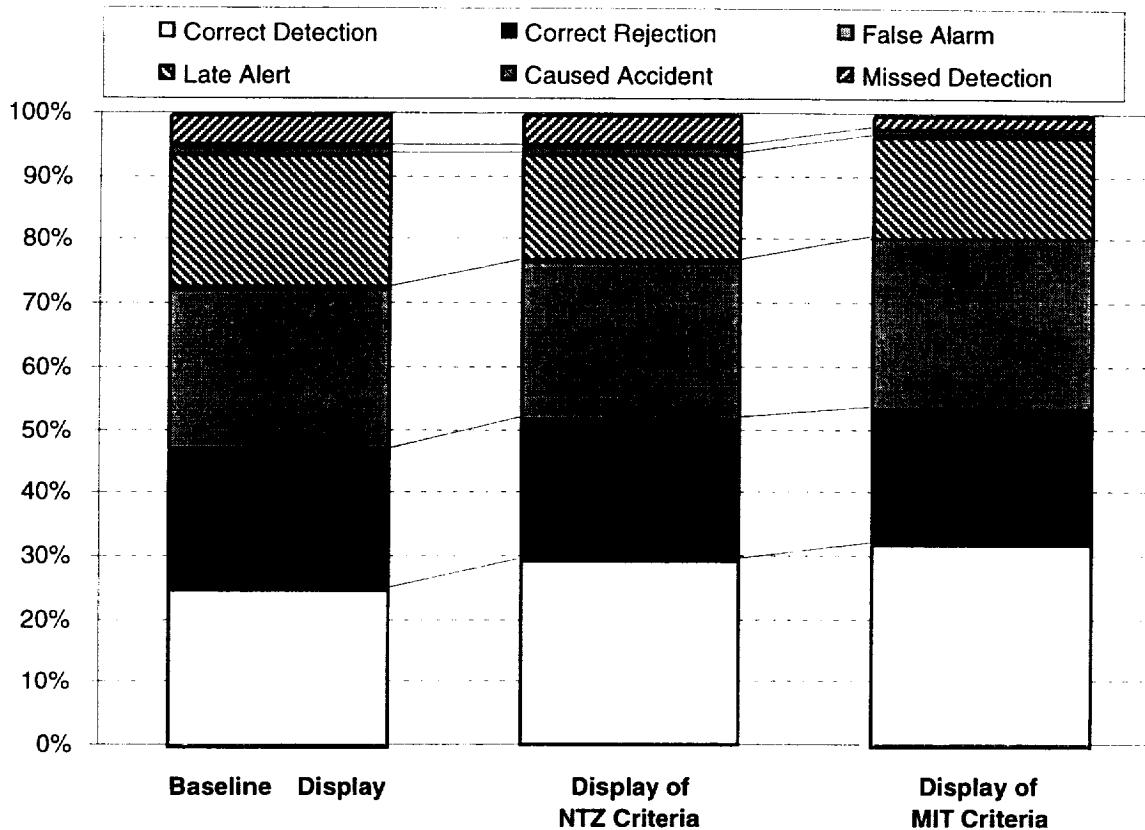


Figure 5.12 Timing of the Subjects Reactions in the High Convergence, Hazardous Blunder Scenario, By Display



**Figure 5.13 Performance of Subjects' Reactions with the Different Displays
(n = 144 For Each Display)**

This shift in the subjects' reactions with the display of the MIT criteria caused a slight improvement in performance. As shown in Figure 5.13, subjects' reactions when the MIT criteria display represented more Correct Detections of an impending collision, compared to their reactions with the baseline display. However, this difference was only significant at the $p < 0.10$ level.

It should be noted that this gain in performance did not result in consistently acceptable performance. A non-negligible number of Missed Detections and Late Alerts occurred with each of the displays.

5.4.3 Effects of the Automatic Alerts

In two-thirds of the runs, automatic alerts were given, based on the NTZ criteria in one third of the runs and on the MIT criteria in the other third. In the remaining runs, no automatic alerts were given. Subjects were not required to follow the automatic alerts. This section examines the effects of automatic alerts on the characteristics of the subjects' reactions.

Paired-comparison statistical tests found several statistically significant differences in the subjects' reactions between the different automatic alert conditions. In the high convergence scenarios, when the MIT criteria generates alerts earlier, the subjects' reactions were earlier when automatic alerts based on the MIT criteria were shown than when automatic alerts based on the NTZ were shown. ($p < 0.01$ in the high convergence, hazardous blunder case; $p < 0.05$ in the high convergence, less-hazardous blunder case). In the low convergence scenarios, when the NTZ criteria generates alerts earlier, subjects' reactions tended to be earlier.

Several trends were noted which may indicate subjects felt the automatic alerts based on the MIT criteria tended to be too early. First, subjects tended to react later after an automatic alert based on the MIT criteria than an automatic alert based on the NTZ criteria; this difference, however, has only a statistical significance of $p < 0.10$. Second, subjects' subjective ratings of the automatic alerts' timeliness tended to describe the MIT criteria alerts as being early more often than the NTZ alerts; this trend also only has a statistical significance of $p < 0.10$. Third, subject ratings of certainty in their decisions were higher when they were not reacting after an MIT-criteria automatic alert ($p < 0.05$ when compared with the cases where no automatic alerts were given).

During the debriefing, subjects were asked for free responses to the question "How did the (automatic) alerts affect your decisions?" The responses were categorized and the number of responses in each category are shown in Figure 5.13. These responses indicate a tendency for the subjects to perceive their decision-making process to be affected by the availability of automatic alerts in three ways:

- The automatic alerts may have been used as additional input to the subjects' reasoning, as shown by subject responses 'Additional Consideration in a Marginal Decision' (6 of 12 responses) and 'Made Me Reconsider Safe Situation' (2 of 12 responses).
- The automatic alerts may have served as a cue for the subjects to evaluate the situation, as shown by the subject responses 'Didn't Scan As Often / Attention Getter' (4 of 12 responses).
- The automatic alerts may have given the subjects greater trust in their decisions when an automatic alert coincided with their reaction, as shown by 'More Trust or Confidence if (Automatic) Alert Also Given' (2 of 12 responses).

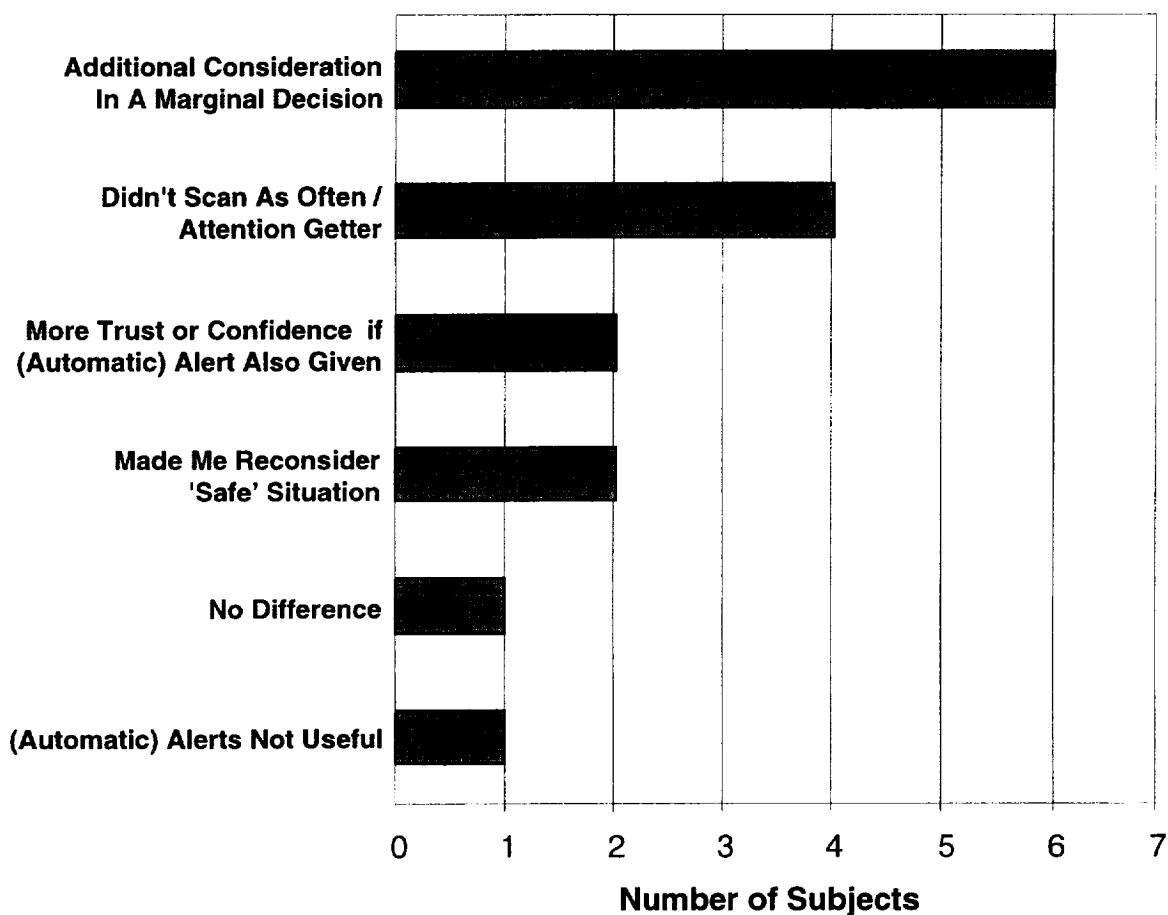


Figure 5.14 Number of Free Responses in Each Category to the Question "How Did The (Automatic) Alerts Change Your Decisions?" (12 Subjects)

5.4.4 Combined Display & Alert Effects

When no automatic alerts were given, the subject's reactions appeared to be strongly correlated with any criteria shown explicitly on the display, as shown by the time difference between the subjects' reactions and the times when each of the alert criteria would have triggered. The mean values of these differences are shown in Figure 5.15. The average difference between the subject's response time and the time the NTZ criteria triggered is significantly different when the NTZ criteria is shown compared to when the baseline display is shown ($p < 0.01$). A similar effect is found for the MIT alert criteria, with a statistically significant difference between subject's reactions with the baseline display available and with the display of the MIT criteria ($p < 0.05$).

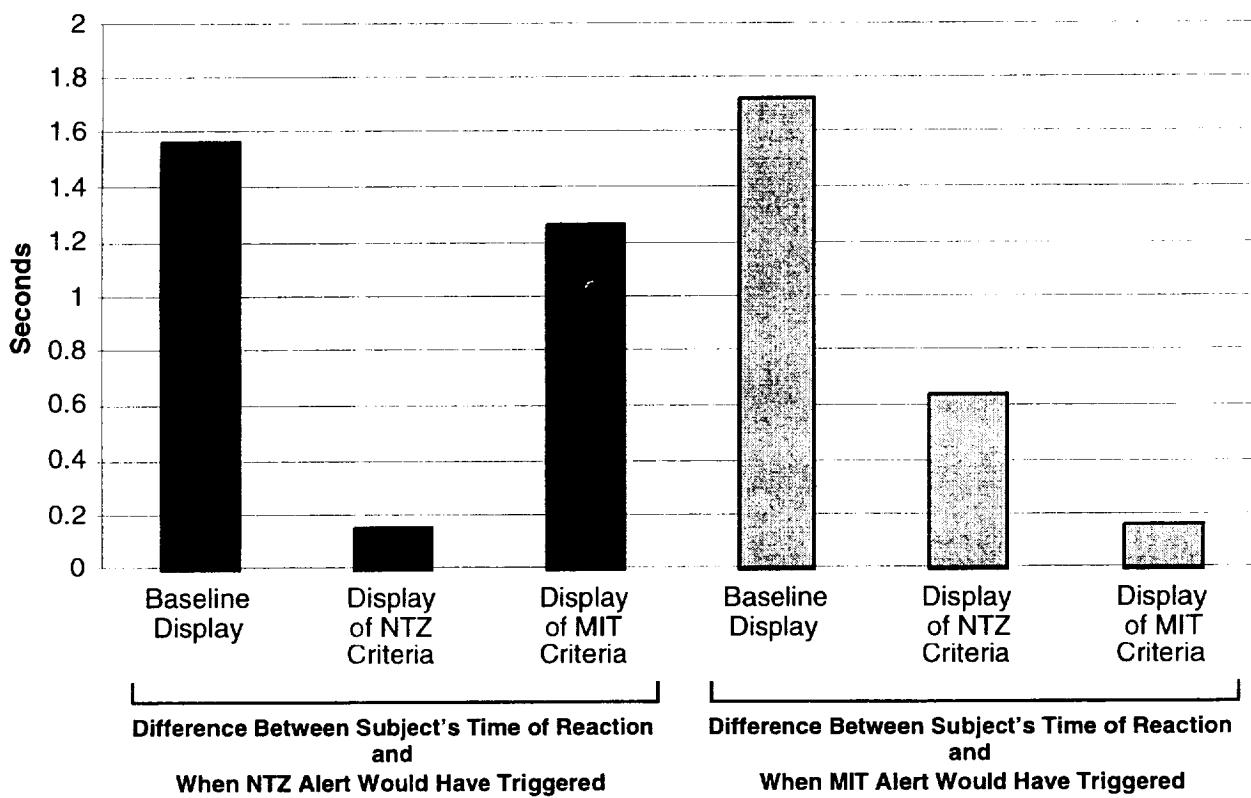
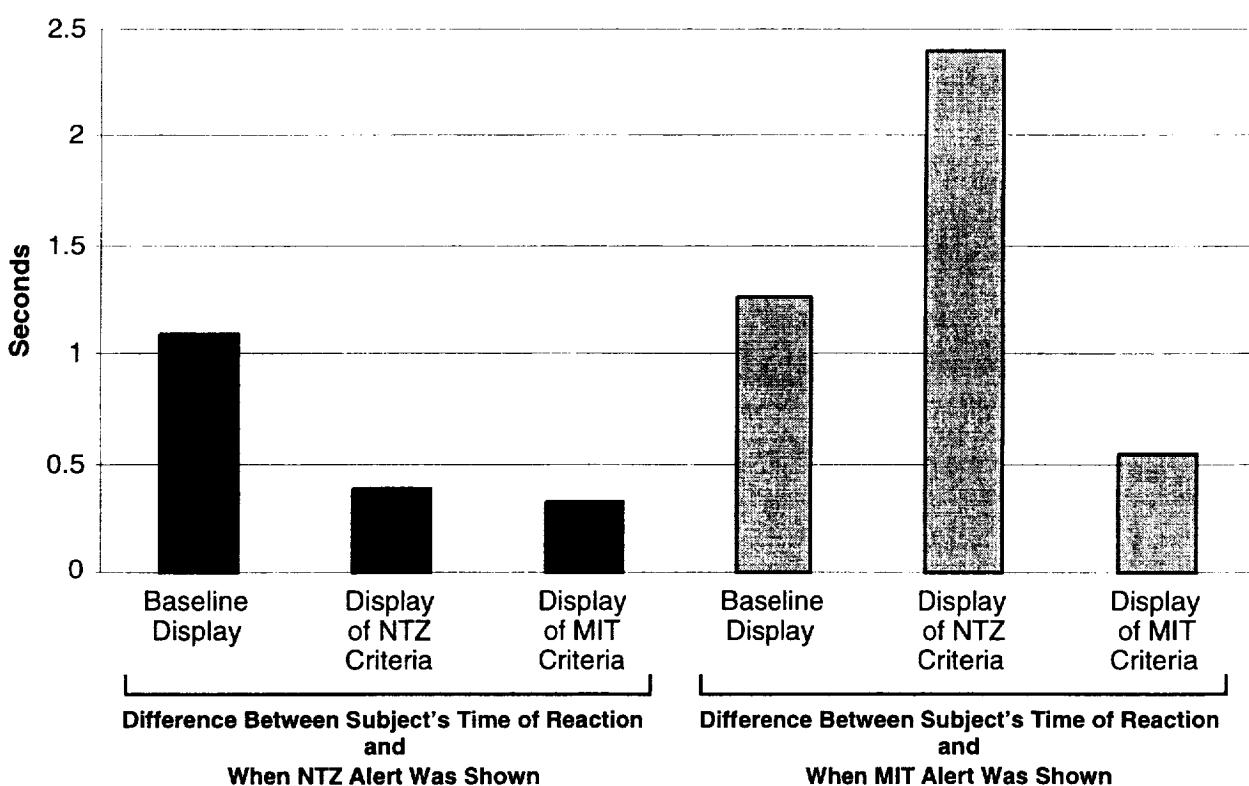


Figure 5.15 Differences in Time Between Subjects' Reactions and Both Alerting Criteria, When Alerts Are Not Given, By Display
(n between 30 & 36 in each condition)

In general, consonance between the criteria on the subjects' displays and the automatic alerts lowered the difference in time between the subjects' reactions and the time when each type of automatic alerts were given, as shown in Figure 5.16. Responses to automatic alerts based on the MIT criteria were the quickest when the MIT criteria was explicitly shown on the display. In contrast, subjects' reactions varied the most from the time of the MIT criteria based automatic alerts when the dissonant NTZ criteria was explicitly displayed. However, because subjects' reactions to automatic alerts based on the MIT criteria were variable, statistical significance of these trends can not be proven.

Subjects' reactions were significantly closer to automatic alerts based on the NTZ criteria when either alert criteria was explicitly shown on the traffic display. The mean difference in time between the subjects' reactions and the time of NTZ-based automatic alerts drops significantly from the runs with the baseline display, compared to the display of the NTZ criteria and of the MIT criteria ($p < 0.01$ & $p < 0.05$, respectively).



**Figure 5.16 Average Different in Time Between When the Subject Reacted and an Automatic Alert Was Triggered by Each Type of Alert Criteria, By Display
(n between 30 & 36 in each condition)**

These differences in subjects' reactions had effects on their final performance. For example, the frequency of Late Alerts by the subjects in the high convergence rate, hazardous blunder scenario is shown in Figure 5.17. Fewer Late Alerts are found when the MIT criteria is displayed to the pilots, compared to other two displays. This effect was the strongest when the MIT criteria was explicitly displayed and automatic alerts based on the MIT criteria were given. However, given the low number of simples in each condition, this trend was only significant to the $p < 0.05$ level between display cases in the conditions with NTZ-based automatic alerts.

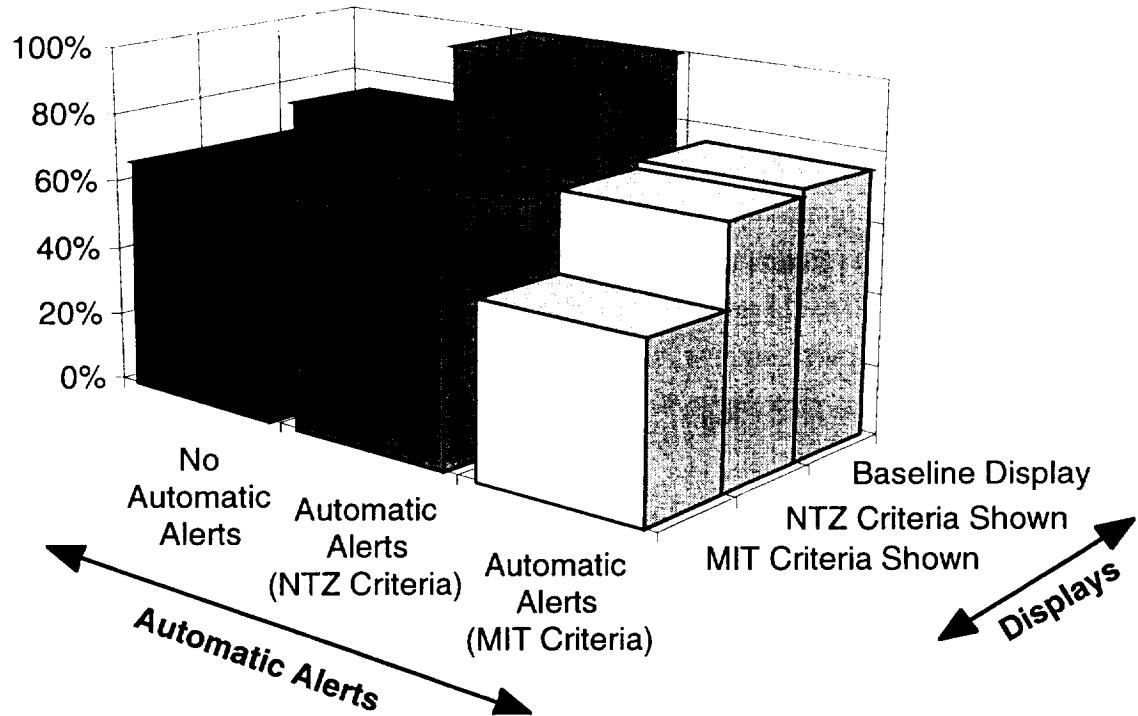


Figure 5.17 Percentage of 'Late Alerts' Given in the High Convergence Rate, Hazardous Blunder ($n = 12$ in each condition)

5.5 Summary of Results and Discussion

This experiment tested the effects of explicitly displaying alert criteria on the traffic display. In cases where automatic alerts were given, the display of alert criteria was purposefully set to generate consonance and dissonance with the automatic alerts in order to test the effects of the display on subject acceptance of, and agreement with, automatic alerts. Two alert criteria were used: a simple, static NTZ criteria which is consistent with subject reactions in Experiment #2, and

a high performance -- but more complex and dynamic -- prototype alert criteria developed at MIT by Carpenter and Kuchar (1996).

The subjects' reactions tended to be more consistent with the alert criteria explicitly presented on the traffic display. For example, the subjects' reactions were closer in time to when the displayed criteria would have triggered.

Automatically generated alerts also affected the subjects' reactions. Subjects' reactions were earlier in cases where automatic alerts were generated before the subjects would normally react. This effect was supported by subjective comments, which suggested an ability of automatic alerts to sway subjects' decisions in marginal cases and to bring a hazardous situation to the subjects' attention.

Specific cases were examined where the criteria explicitly displayed to the subject was also the basis for an automatic alert (consonance) and where the criteria on the display differed from the basis for the automatic alert (dissonance). In cases with consonance between the displayed alert criteria and the automatic alert, the subjects tended to follow the automatic alert more closely. In the case with dissonance between automatic alerts based on the higher-performance MIT criteria and the display of the subject-preferred NTZ criteria, however, subject reactions varied the most from the automatic alerts.

This results provide insight into the relative effects of automatic alerts and the explicit display of alert criteria, and highlight the importance of consonance between the displays and the automatic alerts. Several practical considerations for the task of closely spaced parallel approaches require further study, however. Although the display of the MIT criteria appears to have shifted the subjects' reactions towards more effective alerting strategies, the display of the MIT criteria did not completely meet the ultimate objective of enabling the subjects to consistently use strategies good enough to ensure collision avoidance. Despite a trend of fewer Late Alerts and False Alarms when the MIT criteria was shown, these rates remained high with few highly statistically significant differences in final performance between conditions. These high rates may be an artifact of the difficult scenarios in the experiment, or it may indicate the alerting task is difficult and sensitive to variability in the subjects' reactions.

In addition, the display of the MIT criteria -- or a similar criteria -- may not be the final or best display to provide to pilots, as its often large changes in size and shape may provide a graphical indication too noisy and distracting for a display central to the pilot's scan during final approach.

6. Pilot Conformance to Automatically Generated Commands

In the preceding chapters, a series of flight simulator experiments examined the collision avoidance task during closely spaced parallel approaches. Experiment #1 found frequent instances of pilot non-conformance. For example, when an alerting system did command avoidance maneuvers, the pilots did not follow them approximately 40% of the time. Experiments #2 and #3 examined possible causes of, and solutions to, pilot non-conformance during this task.

These results raise broader issues about pilot interaction with executive alerting systems. This chapter examines pilot interaction with alerting systems, both with and without pilot conformance. Factors which may contribute to non-conformance are examined. The resulting implications of non-conformance on pilot workload and on the resolution to the hazard executed by pilots are discussed. Finally, considerations for alerting system design are given.

6.1 Pilot Interaction With the Alerting System, With and Without Conformance

Alerting systems with executive roles are designed with the implicit assumption that pilots will execute the alerting system commands quickly and precisely. For example, in the executive decision-making alerting system shown in Figure 6.1, the alerting system is assumed to take over completely all of the components of the alerting task except for control actuation. Conversely, the pilots are assumed to perform none of the components except for control actuation.

Non-conformance may be intentional or unintentional. In cases of unintentional non-conformance, pilots are attempting to use the alerting system in the manner intended by its designers, but can not because of conflicting concerns, a lack of awareness of the alerting system's commands, or confusion about what is being commanded.

This thesis has focused on intentional non-conformance to the alerting system's commands. In cases of intentional non-conformance, pilots understand the alerting system's commands, but instead elect to perform some elements of the alerting task, and execute a resolution to the hazard which may or may not resemble that commanded by the alerting system. In Experiment #1, for example, the pilots demonstrated familiarity with the alerting system in the cases where they followed its commanded avoidance maneuvers. Therefore, for the pilots to execute a different avoidance maneuver suggests they were evaluating the situation themselves, and sometimes not following the alerting system's alerts and commands exactly.

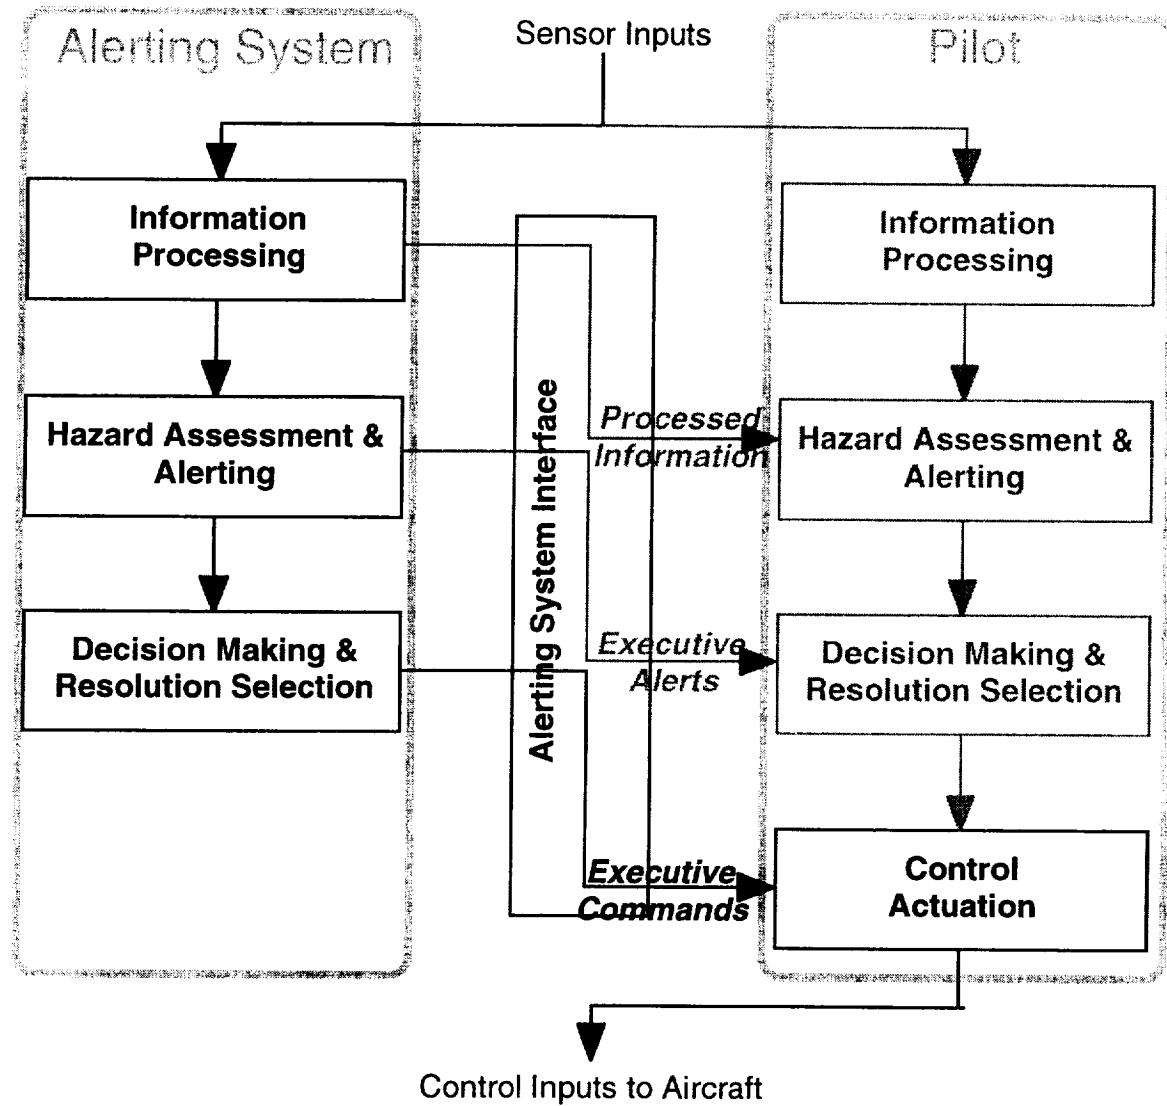


Figure 6.1 Expected Development of Control Inputs to the Aircraft with an Executive, Decision-Making Alerting System Available

For pilots to perform any of the alerting task sub-functions in parallel with the alerting system effectively changes the role of the alerting system and of the pilots' task. As modeled in Figure 6.2, the alerting system outputs to the pilots are used as information sources instead of executive commands. The pilots may also consider information not used by the alerting system. The pilots then, through reconciliation of their own decisions and the commands displayed by the alerting system, decide on a hazard resolution.

This reconciliation can happen for any components of the alerting task. For example, in Experiment #1, pilots did not conform to the alerting component of the alerting system by executing an early avoidance maneuver in 17% of the approaches, before an automatic alert. In

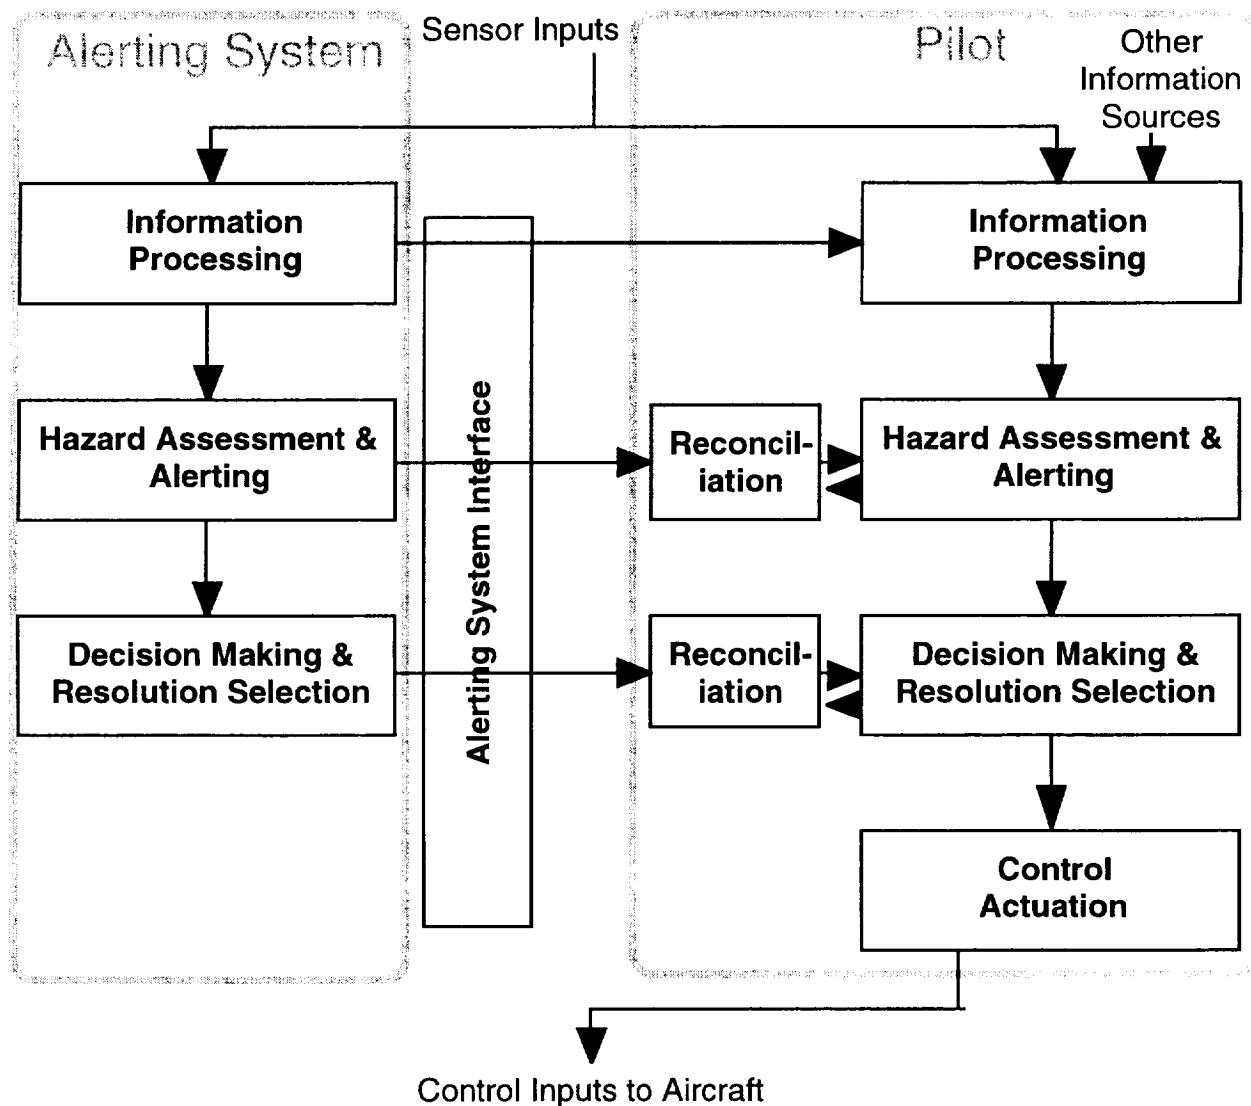


Figure 6.2 Actual Development of Control Inputs to Aircraft When Pilot Does Not Conform

40% of the remaining cases, pilots reacted immediately after the alerting system's alert, but then executed a different avoidance maneuver than commanded, thereby disregarding the decision-making portion of the commands.

Several different methods of reconciliation may be used. At its simplest, the pilots may choose to ignore the alerting system entirely. With other forms of reconciliation, the pilots may use the information from the alerting system in several ways. For example, subject responses in Experiment #3 identified a benefit of having the alerting system act as a trigger to the pilot to evaluate the situation (4 out of 12 responses), and a willingness to factor an automatic alerting into their own evaluation of the problem as an additional piece of information (8 out of 12 responses). Appendix F compares several possible methods of reconciliation in greater detail.

The frequency with which pilots perform these re-evaluation and reconciliation processes may be higher than the non-conformance rate measured in the experiments, which noted only when the pilots' reactions differed from the alerting system's commands. When the pilots did follow the alerting system commands, it is unknown whether they were trusting them without any evaluation, or whether they were taking on the extra workload of evaluation and reconciliation, and then accepting the alerting system's commands.

6.2 Factors Contributing to Pilot Non-Conformance

Pilot non-conformance to an alerting system's commands implies the pilots perceive a need to verify the alerting system's commands. This perception requires pilots to evaluate the situation and reconcile their decisions with the alerting system's commands. Depending on the results of this reconciliation, one of two possible outcomes can occur: pilots may decide to follow the alerting system's commands, or they may decide to execute a different resolution (or no resolution) to the hazard. Either outcome may add a delay to the pilots' responses due to the evaluation and reconciliation processes.

This section will discuss two factors contributing to pilot non-conformance to alerting system commands. First, the pilots' perception that the alerting system's commands need confirmation will be discussed. Then, mismatches between the pilots' decisions and the alerting system's commands will be examined.

6.2.1 Pilot Desire to Confirm the Alerting System's Commands

This section will discuss the types of concerns involved in the pilots' perceived need to confirm the alerting system's command. These concerns include:

- The pilot may be concerned that the alerting system will fail to act as it should.
- The pilot may feel the alerting system can not consider relevant information or has different objectives.
- The pilot may place greater confidence in their own decisions than in the alerting system's.

Pilots' confirmation of the alerting system's commands may stem from a concern that the alerting system will fail to act as it should. These failures may be of two types, each of which have different implications.

The first type of failure occurs if the alerting system either fails to identify a problem, or does not command sufficient action to remedy a problem. In the case of a collision avoidance system, for example, alerting system failures of this type include:

- Failure of a collision alerting system to identify a traffic conflict.
- Generation of a collision alert too late for an effective avoidance maneuver.
- Generation of an avoidance maneuver command which worsens the traffic hazard.
- Generation of an avoidance maneuver command which, while lessening the traffic hazard, does not create adequate aircraft separation.

The decision of pilots to monitor for these types of failures may have several causes. First, the direct effects of these failures can have very high costs; in the case of a collision avoidance system, this type of failure can have catastrophic results. Second, it may be difficult for pilots to develop confidence in the alerting system. Some alerting systems are designed to monitor for rare events. In this case, pilots will not see the alerting system perform correctly in enough instances to build up trust in the system.

Pilots may have exhibited concerns about this type of failure in Experiment #1. For example, a large number of Early Avoidance Maneuvers were generated before the alerting system commanded a maneuver (17% of the approaches). This may be an indication the pilots were more conservative than the alerting system and were concerned, in those instances, that the alerting system would not act early enough.

Pilot concern about this type of failure has several implications. First, if the pilots are not confident that the alerting system will generate an alert when required, they may feel compelled to assess the situation regularly independent of the alerting system. Second, if the pilots feel the commanded resolution to the hazard is insufficient, they may feel compelled to make their own

decisions about a resolution to the hazard, or they may execute a more severe version of the commanded resolution.

The second type of alerting system failure occurs when the alerting system generates unnecessary or overly conservative commands. In the case of a collision avoidance system, for example, these types of failures would include:

- Displaying a collision alert when not warranted.
- Commanding an avoidance maneuver which is more severe than required to maintain adequate aircraft separation.

When the alerting system is designed to prevent catastrophic events, variance in the sensor measurements and unpredictability in the system dynamics requires its reasoning to be conservative. While a conservative design helps ensure prompt, adequate reactions to dangerous situations, it also increases the frequency of false alarms and excessive commands from the alerting system. For example, in order to provide adequate detection of potential conflicts, TCAS II logic is estimated to generate false alarms in 10% of normal approaches to parallel runways separated by 1800 feet. (Folmar, Szebrat & Toma, 1994)

Although the alerting system is performing to specifications, false alarms may appear to the pilot as failures of the system. For most tasks, the resultant cost of this type of alerting system failure is comparatively low in any single event. However, this type of failures can have indirect, cumulative effects. First, they may degrade the pilots' trust in all information presented by the alerting system by making the alerting system's functioning appear spurious and unreasonable. By appearing to have poor performance at one element of the task, the alerting system becomes susceptible to suspicion that it will not perform well at other elements of the task and will be susceptible to failures of both types.

Second, past experience with second type failures reduces the pilots' future confidence that an automatically generated decision is not also a second type failure. For example, a purely rational assessment of an automatic alarm's validity based solely on historical precedent would estimate the likelihood that an alert is valid -- and not a false alarm -- to be the frequency of correct alarms experienced in the past. This confidence estimate is low when many false alarms have been experienced in the past compared to the number of valid alerts. In cases where the hazard occurs rarely, a modest number of false alarms may significantly lower confidence in future alerts.

Pilots can not be assumed to exactly keep track of the false alarm rate and studies have shown humans do not use purely rational assessments of relative probabilities. (Kerstholt et al, 1996) The general effect that false alarms can lower pilot confidence in automatic alerts, however, is supported by surveys of pilot use of both TCAS II and the Ground Proximity Warning System

(GPWS) for air transport aircraft. (DeCelles, 1991; Ciemier et al, 1993) False alarms were indicated in both cases as the primary factor inhibiting pilots' immediate reactions to automatic alerts.

A concern that an alerting system may be overly conservative has two implications. First, this concern can drive pilots to delay their response to automatic alerts and commands while they confirm their veracity. Second, if pilots do not have complete understanding of the situation, they may erroneously judge automatic commands to be unnecessary, and the expected benefit of the alerting system will not be achieved.

The second factor in pilots' desire to confirm alerting system commands is a perception that, while the alerting system is functioning to its specifications, these specifications do not include knowledge of all information or have the same objectives as the pilots.

For example, pilots indicated in a survey that they sometimes do not follow TCAS commands -- or turn them off -- in conditions where they visual contact with the other aircraft or have knowledge of the other aircraft's intentions through ATC communications. (Ciemier et al, 1993)

Similar results were found in Experiment #1. When a vertical avoidance maneuver was commanded by the alerting system, pilots often used horizontal avoidance maneuvers instead of solely conforming to the vertical avoidance maneuver commanded by TCAS. This may indicate pilots felt they were including more information in their own decisions than the alerting system.

The third factor in pilots' desire to confirm an alerting system's commands is the relative confidence they place in the alerting system compared to their own decisions.

Just as pilots develop confidence in the alerting system, they also associate a confidence with their own reasoning. Like their confidence in the alerting system, this value can be based both upon previous experience and upon criteria which are harder to assess and anticipate. In addition, this self-assessed level of confidence can be dynamic, varying with the pilots' perceptions of the situation and of their ability to make decisions given their immediate workload, knowledge, and situation awareness. Riley (1996) describes pilot self-confidence in decision making skills as impacted by the perceived workload, skill, risk, task complexity and past success at the task. For example, a pilot may not feel confident in high workload conditions at a task which is easily achievable in low workload conditions.

Four general combinations of pilots' confidence in their own decisions and in the alerting system's commands are shown schematically in Figure 6.3. When pilots have a high confidence in their own reasoning and a low confidence in the alerting system's reasoning, they are more likely to act upon their own reasoning and to confirm automatic commands. With a higher relative confidence in the alerting system, pilots will feel less of a need for confirmation of automatic commands.

		Pilot Confidence in Their Own Reasoning	
		High	Low
Pilot Confidence in Automatic Commands	High	Indeterminate	Pilot Will Probably Follow Automatic Commands
	Low	Pilot Will Probably Confirm or Ignore Automatic Commands	Indeterminate

Figure 6.3 Pilot's Actions Based Upon Their Confidence in the Alerting system and Their Own Reasoning

When pilots' self-confidence and confidence in the alerting system are comparable, the final outcome can not be predicted and may be more sensitive to specific features of the immediate situation.

Because the pilots' self-confidence in their decisions may vary with a variety of conditions, which 'quadrant' will be applicable at any given time can not be pre-determined. However, Figure 6.3 helps illustrate both the worst-case scenarios and possible changes in conformance to alerting system commands with features of a specific situation and with changes in pilots' self-confidence through experience and training.

In Experiment #1, pilots were found to conform more often when given the traffic display with the least information about the relative position of the other aircraft. This may be an example of a case where pilots were more willing to rely on the alerting system when they are given too little information to feel confident in deciding upon and executing different resolutions than commanded by the alerting system.

The experiment results focused on subjects' immediate perceptions of the alerting system. These perceptions may have cumulative effects and may add to the longer-term development of pilot 'trust' in the alerting system. The event-driven descriptions of pilots checking for alerting system failures and pilots' confidence in their own reasoning have analogies with the seven

conceptual causes of trust given by Sheridan (1992), paraphrased in Table 6.1. The first five factors pertain to conditions that may reduce the pilots' confidence in the alerting system's decisions. The sixth factor pertains to pilots' need to rely on the alerting system. The seventh pertains to the unintentional non-conformance that may result from pilots being confused about the alerting system commands.

Table 6.1 Comparing Pilot's Low Confidence in Alerting Systems with the 7 Causes of Trust Identified by Sheridan (1992)

<p>Reliability <i>"Repeated, Consistent Function"</i></p> <p>If the alerting system reliably provides correct alerts, then it is assumed the pilots will trust it more. However, if the alerting system reliably provides false alarms or appears to act differently in different conditions, then the pilots may give less credence to its commands.</p>
<p>Robustness <i>"Ability to Perform Under a Variety of Circumstances"</i></p> <p>If the alerting system is designed for frequent events, then the pilots can ascertain its robustness directly. However, if the alerting system is designed for rare events, then the pilots rarely receive direct indications of proper functioning under a variety of conditions.</p>
<p>Understandability <i>"Human Can Form a Mental Model of the System and Predict Future Behavior"</i></p> <p>Depending on the complexity of the underlying logic of the alerting system, training, and the amount of information displayed to the pilots, they may or may not fully understand the reasoning behind the command.</p>
<p>Explication of Intention <i>"System Explicitly Displays or Says It Will Act In a Particular Way"</i></p> <p>Executive-type alerting systems do not necessarily display the underlying criteria determining the automatic commands.</p>
<p>Usefulness <i>"Able to Respond in a Useful Way"</i></p> <p>The pilots may, or may not, believe the alerting system is considering all relevant goals or information in forming its commands.</p>
<p>Dependence <i>"Dependence on the System"</i></p> <p>Depending on their self-confidence in their own reasoning, the pilots may be forced into a dependence on the alerting system, or may have established sufficient trust to depend on it.</p>
<p>Familiarity <i>"Naturalness of Displays and Controls"</i></p> <p>The alerting system must be designed so that the meaning of the alert or command is clear, and so that the command is easily understood.</p>

6.2.2 Differences in Strategies Used by Alerting System and Pilots

If the pilots do not have confidence in the alerting system, they may attempt to confirm its alerts and commands. This confirmation process alone can cause a delay in the pilots' responses.

If the pilots' assessments do not agree with the alerting system's commands, they may additionally execute different resolutions to the hazard. This additional component of non-conformance was found in Experiment #1, when the pilots executed avoidance maneuvers that did not meet the minimum criteria of the alerting system commands in 40% of the cases.

Experiment #2 identified differences between subjects' decisions and automatic commands for the task of collision avoidance during closely spaced parallel approaches. Subjects were found to generate alerts consistent with a simple alerting criteria based on position, even when shown displays of convergence rate and trend. This alerting criteria differs from the more complex logic proposed for alerting systems for closely-spaced parallel approaches, generating the potential for disagreements between the alerting system and pilots.

Experiment #3 found the subjects' reactions were shifted by the explicit display of alert criteria. When the display of alert criteria was consonant with the alerting system's logic, subjects' reactions followed automatic alerts more closely. However, when the display supported different alerting criteria than that used by the alerting system, the subjects' reactions were more varied.

These experiment results illustrate how pilots may prefer simple, unambiguous strategies based on immediately available information. These strategies may vary from the increasingly sophisticated logic being developed for alerting systems. A resulting mismatch may contribute to non-conformance.

6.2.3 Effects of Mismatches on the Pilot's Confidence in the Alerting System

The previous sections discussed pilot confidence in automatic alerting systems and mismatches between the pilots' decisions and the alerting system's commands. These two effects may be strongly related. It should be noted that pilot confidence in the alerting system is driven by perceptions. These perceptions can be formed by many factors, some of which may be beyond the alerting system designer's control, such as preconceptions about the alerting system in particular and automation in general.

Mismatches between the pilot's decisions and the alerting systems commands may have a dominant effect in the pilot's perception that the system is not trustworthy. This perception may be accurate, or it may be flawed. For example, in Experiment #3, subjects tended to rate the

automatic alerts as early, and did not follow them quickly -- even when the automatic alerts were correct and immediate action was required.

These perceptions can be influenced by both present information and past experience. The experiments in this thesis focused on the effects of the displays immediately before the subjects. The pilots' past experience with the system may also affect their confidence. Therefore, any mismatches can have a lingering effect on the pilots' decisions to conform, and it can not be assumed that pilot confidence in a high performance alerting system will be increased through pilot experience with it, should the pilots do not always understand or agree with its commands.

6.3 Effects of Non-Conformance on Alerting System Benefits

Implementation of an executive alerting system is typically expected to increase system performance at some metric, such as an increased ability to resolve traffic conflicts, while eliminating the need for pilots to perform the alerting and decision-making sub-tasks. When pilots instead confirm the alerting system's alerts and decisions, they are effectively changing the role of the alerting system from being executive to being supportive. In doing so, the anticipated benefits of the alerting system may not be fully realized. This section will briefly examine the general effects of non-conformance on the pilot's task and on the pilot's ultimate ability to resolve the hazard.

6.3.1 Implications of Non-Conformance on Pilot Workload

When an alerting system is not available, pilots are expected to perform all components of the alerting task. Some of these components may require vigilance by the pilots during nominal operations, such as continuously checking the need for an alert. When a hazard is projected to occur, the pilots additionally assume the workload of decision-making and control actuation.

By giving an executive role to the alerting system, it is expected that the pilots will be relieved of responsibility for some components of the alerting task. For example, with the executive decision making alerting system shown in Figure 6.1, it is expected that the responsibility for the alerting and decision making components of the alerting task are given to the alerting system; the pilots are responsible only for control actuation. Pilots are sometimes presented with information to allow monitoring of the alerting system, but this monitoring is often assumed to be a passive, low-workload task without any additional processing required by the pilots.

Non-conformance to the alerting system's commands implies that the pilots are still executing some or all of the components of the alerting task. If the alerting system presents sufficient supporting information to make verifying its commands easy, then this workload may be small. However, if the alerting system's commands are difficult to understand, the reconciliation and decision-making tasks may be intensive or the pilots may choose to ignore the alerting system entirely.

This predicted increase in workload may be related to studies finding a higher-than-expected workload at passive monitoring tasks. (Parasuraman et al, 1996; Wiener & Curry, 1980) Additional, unexpected higher-workload tasks may be performed by pilots in conjunction with simply monitoring an alerting system's interface. The description of the pilots' task as including unanticipated cognitive analysis and reconciliation may explain the high workload found by these studies.

6.3.2 Effects of Non-Conformance on System Performance

If the pilot does not follow the alerting system immediately and/or does not execute its commands, the resultant system behavior can no longer be described by the pre-determined functioning of the alerting system and the performance of the system can be affected.

Unlike the known logic underlying the functioning of an executive alerting system, the algorithms pilots will use to formulate their own decisions and to reconcile their decisions with the alerting system's commands can not be predicted with certainty. Involvement of pilots in the decision making removes the ability to analyze the system behavior with the same degree of certainty. This variability may limit the extent to which the performance of the combined pilot-alerting system can be projected during design and certification.

Two general effects on system behavior are possible. First, pilots' separate execution of the sub-task and reconciliation with the alerting system's decisions may add a significant, unexpected time delay to the pilots' responses. This delay may involve more time than is usually considered in the design of alerting system, which looks at basic response time to displayed commands. By evaluating the situation, pilots may add a cognitive component to the response time which may be substantial, which may include purposefully waiting for more information, and which may have a significant effect on performance. This effect was noted in Experiment #3. Subjects tended to take longer to react after alerts generated by a different alert criteria than that apparently preferred by the subjects, and therefore triggered alerts which were too late for an effective avoidance maneuver approximately 70% of the time in the most hazardous scenario.

Second, in the presence of a mismatch, the behavior of the system will tend to be skewed towards that the pilots would command if unassisted by an alerting system. This may be a positive effect in conditions where the pilots have more information or a better strategy than the alerting system. However, in cases where the pilots' decisions are deficient, the overall system performance may be degraded. The extent to which pilots' actions differ from the automatic commands depends on the method of reconciliation used and the pilots' trust in the alerting system's commands. In Experiment #3, this difference was also found to be related to the amount the subjects' displays supported the strategies used by the alerting system.

6.4 The Impact of Alerting System Design on Pilot Conformance

Based on the discussion in the previous sections, pilot non-conformance is thought to be influenced by three conditions:

- *Low Pilot Confidence in the Alerting System* This confidence was given as the impetus for pilots to perform any additional reasoning beyond that required for the simple execution of the alerting system's commands. Pilot confidence may be influenced by many factors, including a perception that the alerting system is prone to failures or does not consider all relevant information.
- *Mismatches Between the Pilots' Reasoning and the Alerting System's Commands* As alerting systems become increasingly sophisticated, they can generate commands which are based on different, more complex criteria than involved in the pilots' reasoning. As such, pilots may feel that the commands are erroneous. If the pilots place more trust in their own reasoning, they may not conform to the alerting system.
- *Difficulty for the Pilots in Reconciling Their Decisions and the Alerting System's Commands* Pilots may attempt to compare their decisions to the alerting system's commands. If the basis for the alerting system's commands is available to the pilots, they may be more convinced to conform and they may change their own reasoning about the task. Conversely, if the pilots can not perceive a rationale for the alerting system commands, then they may ignore them or judge them to be spurious.

In cases where pilot non-conformance may have detrimental effects, two possible methods of promoting pilot conformance can be envisioned. First, the alerting system's commands may be made mandatory for the pilots to follow, and any displayed information which might give the pilots confidence in non-conforming actions may be removed. Although this method promotes pilot conformance, it also raises several issues. In reducing the role of the pilot to an un-informed control actuator, the anticipated benefits of having a pilot in the loop -- such as having a flexible

decision-maker capable of correcting any failures or un-needed commands of the alerting system -- are lost. This role of the alerting system may also generate strong pilot opinions and have difficulty being accepted. Finally, blocking off all relevant information may be difficult; pilots may use information from other sources as a basis for non-conformance.

The second method encourages informed decisions by the pilots. In situations where the alerting system's commands are valid, this method promotes pilot conformance, while maintaining the benefits of a pilot in the loop in situations where pilots have better reasoning. As such, this method has two design objectives:

- To reduce mismatches between the pilots' decisions and the alerting system's commands, explicitly present the *synthesized information* implicit in the alerting system's algorithms.
- To make the task of reconciling the pilots' decisions and alerting system commands easier, explicitly present the *alerting thresholds and decision-making objectives* used by the alerting system's algorithms.

For example, the hazard assessment and alerting function implicitly contains intermediate steps. Given the current state of the system, the future behavior is predicted and the hazard level is calculated -- *synthesized information*. This synthesized information is then evaluated to determine the need for an alert; this determination is performed in alerting systems by comparisons to predetermined *alert thresholds*.

Pilots are not always capable of the computations required to evaluate the synthesized information. For example, the experiments in this thesis found subjects' unassisted alerting decisions appeared to be consistent with algorithms which compared the position of an intruder aircraft to a static alert threshold. This differs from the computations and stored database required by the proto-type higher performance alerting algorithms. By having the alerting system display a calculated, unambiguous assessment of the synthesized information about the hazard, pilots may be encouraged to use more sophisticated algorithms. This may have the effect of reducing mismatches between the pilots and the alerting system.

In order to reduce the pilots' difficulty with reconciliation, the alerting system can additionally display the alert thresholds or criteria used to evaluate the need for an alert. This feature may have two possible benefits. First, it allows pilots to quickly determine the underlying rationale for the alerting system commands. Second, it may allow pilots to incorporate in an informed manner considerations beyond the pre-determined criteria used by the alerting system.

The concept of displaying hazard information and alert thresholds was studied in Experiment #3. Results indicate a shift in the subjects' decisions towards whichever alert criteria was

displayed on the screen, and a greater acceptance of automatic alerts consonant with an explicitly displayed alert criteria.

Experiment #3 focused on supporting the alerting sub-function. When the alerting system is also responsible for deciding on an appropriate resolution to the hazard, a similar display may be required to support the decision-making sub-function. Like the alerting function, the decision-making function can involve several intermediate steps. First, possible resolutions need to be selected and their performance predicted; this step can be too high in computations for pilots to rigorously execute. Then, based on their anticipated performance, one resolution method must be selected; this judgment may be well suited for pilots to execute themselves or to oversee.

In the task of collision avoidance during closely spaced parallel approaches, alerting systems under development are implicitly assuming responsibility for commanding avoidance maneuvers. However, these alerting systems are also attempting to reduce a mismatch between the commanded avoidance maneuver and that preferred by pilots by always commanding the same, pre-set maneuver. This maneuver is generally the turning, climbing maneuver selected by subjects in the experiments in this thesis, suggesting pilots will be less likely to question the commanded maneuver's validity and more likely to conform.

A longer range objective of alerting system design may be to generate sufficient pilot trust that pilots do not feel a need to confirm or ignore the alerting system's commands. This trust may be contributed to by the consistent reduction in mismatches resulting from pilots misunderstanding the underlying logic for automatic commands, and by the pilots' ability to understand and verify the alerting system's logic.

Similar methods have been published to this thesis' concept of giving automatic systems a combined executive role (to command high performance resolutions to hazards) and supportive role (to promote pilot conformance to the commands). The objectives of this supportive role -- to present the pilots with any highly processed information required for the desired solution algorithms, and to display the objective function used by the alerting system in considering this information -- have also been discussed for a variety of systems. (e.g. Sheridan, 1992)

7. Conclusions

7.1 *Collision Avoidance During Closely Spaced Parallel Approaches*

In order to examine issues with non-conformance, this thesis used the task of collision avoidance during closely spaced parallel approaches as a case study. Numerical analysis of possible trajectories illustrated the time-critical nature of this task. The limited range of potentially dangerous relative aircraft positions -- the 'Kill Zone' -- was identified. The Kill Zone provides insight for the development of traffic displays and approach procedures. The Kill Zone also highlights the potentially non-intuitive alerting criteria required for this alerting task.

At the start of this thesis, relatively little information was available about the manner in which a pilot would use a cockpit collision avoidance system during closely spaced parallel approaches without intervention from Air Traffic Control. Experimental results identified both operational implications for closely spaced parallel approaches in Instrument Meteorological Conditions, and broader implications for pilot non-conformance to alerting systems in general.

A collision avoidance system is needed, as indicated both by pilot opinions and by the less frequent rates of loss of aircraft separation achieved in Experiment #1 when an alerting system issued traffic alerts and commanded avoidance maneuvers.

However, the full benefit of an alerting system may be difficult to realize because of pilot non-conformance. In Experiment #1, pilots appeared to intentionally not conform to the minimum specifications of alerting system commands in 40% of the cases. Pilot non-conformance was found to have a dominant and detrimental effect on the rate of loss of aircraft separation.

Displaying more information to the subject-pilots about the other aircraft relative to their approach path had a negligible or negative effect. In Experiment #1, the non-conformance rate was significantly higher when the pilot preferred 'Enhanced' traffic displays were available, possibly indicating these displays tended to give pilots enough confidence to trust their own decisions more than the alerting system's commands. Experiment #2 found the display of intruder aircraft heading and trend did not appear to change the characteristics of the subjects' reactions. These results suggest methods of encouraging pilot conformance at this task must provide the pilot with more information than the current traffic situation alone.

Disagreements between pilots' decisions and alerting system commands may have been a contributing factor to pilot non-conformance. Experiment #2 found subjects' alerts were consistent with a Non-Transgression Zone type criteria which can cause false alarms in less-hazardous situations and late alerts in hazardous situations. Likewise, the subjects selected avoidance

maneuvers which turned away from the intruder in 91% of the cases. These choices differed from the vertical maneuvers commanded by TCAS in Experiment #1, and indicate a consistent tendency to pick the same avoidance maneuvers without necessarily considering the characteristics of a particular traffic situation.

The explicit display of alerting criteria to subjects was found in Experiment #3 to have beneficial effects. The characteristics of the subjects' reactions shifted towards a displayed alert criteria. When the displayed alert criteria was consonant with an automatic alert, the subjects' reactions followed the automatic alert more closely and consistently. These results provide considerations for the design of alerting systems for this task which encourage pilot conformance to higher performance traffic alerts and commanded avoidance maneuvers.

These results may not represent a final solution, however. Even with consonance between the subjects' displays and automatic alerts, subjects tended to rate the higher-performance alerts as early and delayed their responses.

7.2 Pilot Non-Conformance to Alerting System Commands

This thesis has focused on intentional non-conformance to executive alerting system commands. Three experiments examined this issue in a case study. Their results, in addition to identifying specific concerns about the task of collision avoidance during closely spaced parallel approaches, also have broader implications for pilot conformance to alerting system commands and alerting system design.

For pilots to intentionally act differently than commanded by the alerting system implies pilots are assessing the situation in parallel with the alerting system. This effectively shifts the role of the alerting system to be less executive. Pilots may choose to ignore the alerting system commands, or they may attempt to reconcile their decisions with the alerting system's commands.

The impetus for pilots to evaluate the situation and reconcile their decisions with the alerting system may be a lack of confidence in the alerting system. Several factors affecting pilot confidence were discussed. These may include a concern the alerting system will not perform as expected, a perception that the alerting system may not include all relevant information in its decision making, and a higher confidence by the pilots in their own reasoning.

Pilot non-conformance adds variability to the final behavior of the system which may require consideration in the design and evaluation of alerting systems. If pilots decide to evaluate the alerting system's commands, they can add a significant, unexpected time delay to their reaction. Additionally, if there are mismatches between the pilots' decisions and the automatic commands,

pilots may execute a different resolution to the hazard than expected. When the pilots have more insight, more information or better reasoning about the situation, this non-conformance is a positive influence. However, new alerting systems are being implemented which are considered 'better' at avoiding specific hazards. In this case, pilot non-conformance may seriously reduce system performance.

Based on these concepts, considerations to encourage informed pilot conformance have been given for alerting system design. These considerations focus both on making synthesized information available to the pilots to enable the pilots to use higher-performance strategies, and on presenting the decision-making criteria used by the alerting system to the pilots, in order to allow the pilots to understand the basis for automatic commands.

7.3 Recommendations

The high rates of non-conformance found with operational alerting systems and in Experiment #1 have served to illustrate the need to design alerting systems with the objective of promoting pilot conformance. To this end, this thesis has developed considerations for the design of automatic alerting systems which may encourage informed conformance by pilots. These design considerations present a practical emphasis on shaping the strategies used by pilots towards those used by the alerting system, by considering the specific processed information which must be explicitly displayed to pilots in order for the strategies to be used. Because the prototype displays tested in this experiment did not fully resolve problems with inadequate subject reactions, further research into these types of displays is suggested.

These design considerations must be considered not only in the design of the alerting system interface, but also in the design of the alerting system logic. These considerations have focused on communicating to pilots the basis of the underlying algorithms; in order to achieve this, the algorithms must be designed to be communicable.

These design considerations may not always be practical or possible. The executive role of alerting systems has been, in part, motivated by a desire to not increase the amount of information required for pilots to monitor with the introduction of an alerting system and to thereby not increase pilot workload. These design considerations run counter to that objective by proposing the display of additional information. Cases may exist where pilots may not have the time to consider this information or the display of the information may not be possible.

Likewise, these design considerations assume the core elements of the alerting system's algorithms can be displayed to pilots in a manageable, comprehensible manner. As the alerting

tasks become increasingly complex and their performance specifications become increasingly stringent, the desired algorithms may necessarily become too multi-dimensional and too sensitive to variations in interpretation to have their properties displayed.

When these limitations are reached, pilots may not be able to comprehend the reasoning behind the alerting system's functioning and a problem with non-conformance may occur. These effects may therefore represent a limitation on the use of alerting systems as much as a limitation of the design considerations give in this thesis.

References

Andrews, J.W., Senne, K.D. & Koegler, J.C. "IPC Design, Validation and Flight Testing Final Report", M.I.T. Lincoln Laboratory Report FAA-RD-77-150, March 1978

Billings, C. "Situation Awareness Measurements and Analysis: A Commentary" in *Experimental Analysis and Measurement of Situation Awareness*, Ed. Garland, D.J. & Endsley, M.R., November 1995

Cacciabue, P.C. "A Methodology of Human Factors Analysis for Systems Engineering" *Proceedings of the 1993 International Conference on Systems, Man and Cybernetics*, Le Touquet, France, October 1993

Carpenter, C. & Kuchar, J.K. "Probability-Based Collision Alerting Logic for Closely-Spaced Parallel Approaches", to be presented at *AIAA 35th Aerospace Sciences Meeting*, Reno NV, January 1997

Ciemier, C.E., Gambarani, G.P., Sheridan, E.H., Stangle, S.C., Stead, R.P. & Tillotson, D. H. "The Results of the TCAS II Transition Program (TTP)", The ARINC Research Corporation, April 1993

Corjon, A. & Poinsot, T. "Vortex Model to Define Safe Aircraft Separation Distances", *Journal of Aircraft*, Vol. 3, No. 3, May-June 1996

DeCelles, J.L. "The Delayed GPWS Response Syndrome", Aviation Research and Education Foundation, Herndon VA, July 1991.

Endsley, M. "Toward a Theory of Situation Awareness in Dynamic Systems", *Human Factors*, Vol. 37, No. 1, March 1995.

Federal Aviation Administration ARD-300 "Precision Runway Monitor Demonstration Report", FAA Report No. DOT/FAA/RD-91/5

Folmar, V.A., Szebrat, X.P. & Toma, N.E. "An Extension to the Analysis of Traffic Alert and Collision Avoidance System (TCAS) Advisories During Simultaneous Instrument Approaches to Closely Spaced Parallel Runways", MITRE Corp. Document MTR 94W0000056, July 1994

Hersh, M. & Baker, R. "Decision-Making and Optimization", *Proceedings of the 1993 International Conference on Systems, Man and Cybernetics*, Le Touquet, France, October 1993

Hess, R.A. "Structural Model of the Adaptive Human Pilot", *Journal of Guidance and Control*, Vol. 3, No. 5, p. 416-423, September/October 1980

Hogg, R.V. & Ledolter, J. *Applied Statistics for Engineers and Physical Scientists*, Macmillan, New York, 1992

Johannsen, G. "Towards a New Quality of Automation in Complex, Man-Machine Systems", *Automatica*, vol. 28, p. 355-373, 1992

Johnson, E.N. & Hansman, R.J. "Multi-Agent Flight Simulation with Robust Simulation Generation", MIT Aeronautical Systems Laboratory Report, ASL-95-2, January 1995

Kerstholt, J.H., Passenier, P.O., Houttuin, K. & Schuffel, H. "The Effect of a Priori Probability and Complexity on Decision-Making in a Supervisory Control Task", *Human Factors*, 38(1), p. 65-78, 1996

Knaeuper, A. & Rouse, W.B. "A Rule-Based Model of Human Problem-Solving Behavior in Dynamic Environments", *IEEE Transactions on Systems, Man and Cybernetics*, vol. SMC-15, No. 6, November/December 1985.

Kockzo, S. "Co-ordinated Parallel Runway Approaches", NASA Contractor Report CR201661, October 1996

Kuchar, J.K. "Methodology for Alerting-System Performance Evaluation", *Journal of Guidance, Control and Dynamics*, Vol. 19, No. 2, p.438-444, March-April 1996.

Kuchar, J.K. & Hansman, R.J. "A Unified Methodology for the Evaluation of Hazard Alerting Systems", MIT Aeronautical Systems Laboratory Report, ASL-95-1, January 1995

Mårtensson, L. "The Aircraft Crash at Gottröra: Experiences of the Flight Crew", *International Journal of Aviation Psychology*, 5(3), p. 305-326, 1995

McRuer, D.T., Clement, W.F., Thompson, P.M. & Magdaleno, R.E. "Minimum Flying Qualities", Vol. 2, Systems Technologies Inc. Technical Report No. 1235-1, June 1989

Mellone, V.J. & Frank, S.M. "Behavioral Impact of TCAS II on the National Air Traffic Control System", *7th International Symposium on Aviation Psychology*, Columbus OH, April 1993.

Owen, M.R. "The Memphis Precision Runway Monitor Program Instrument Landing System Final Approach Study", M.I.T. Lincoln Laboratory Report No. DOT/FAA/NR-92/11, May 1993

Parasuraman, R., Molloy, R., Mouloua, M. & Hilburn, B. "Monitoring of Automated Systems", in *Automation and Human Performance: Theory and Applications*, Ed. by Parasuraman, R. & Mouloua, M., Lawrence Erlbaum, Mahwah NJ, p. 91-116, 1996

Pawlak, W.S., Brinton, C.R., Crouch, K. & Lancaster, K.M. "A Framework for the Evaluation of Air Traffic Control Complexity", *AIAA Guidance, Navigation and Control Conference*, San Diego CA, July 1996.

Patrick, N.J.M. "Decision-Aiding and Optimization for Vertical Navigation of Long-Haul Aircraft", Ph.D. Thesis, Massachusetts Institute of Technology, August 1996

Radio Technical Committee on Aeronautics (RTCA). "Minimum Performance Specifications for TCAS Airborne Equipment", Document No. RTCA/DO-185, Washington DC, September 1983

Rasmussen, J. *Information Processing and Human-Machine Interaction*, North-Holland, New York, 1986.

Rice, J.A. *Mathematical Statistics and Data Analysis*, Wadsworth & Brooks/Cole, Pacific Grove, CA, 1987

Riley, V. "Operator Reliance on Automation: Theory and Data", in *Automation and Human Performance: Theory and Applications*, Ed. by Parasuraman, R. & Mouloua, M., Lawrence Erlbaum, Mahwah NJ, p. 19-36, 1996

Sarter, N. & Woods, D. "How in the World Did We Ever Get into That Mode? Mode Error and Awareness in Supervisory Control", *Human Factors*, Vol. 37, No. 1, p.5-19, 1995

Shank, E.M. & Hollister, K.M. "Precision Runway Monitor", *The Lincoln Lab Journal*, vol. 7, 1994

Sheridan, T.B. *Telerobotics, Automation and Human Supervisory Control*, MIT Press, Cambridge MA, 1992

Thordesen, M., Militello, L., & Klein, G. "Determining the Decision Requirements of Complex Flight Crew Tasks", *7th International Symposium on Aviation Psychology*, Columbus OH, April 1993.

Toma, N.E. & Massimini, S.V. "An Analysis of Traffic Alert and Collision Avoidance System (TCAS) Advisories During Simultaneous Instrument Approaches", MITRE Working Paper 93W0000222, September 1993.

Wickens, C. "The Tradeoff of Design for Routine and Unexpected Performance: Implications of Situation Awareness", in *Experimental Analysis and Measurement of Situation Awareness*, Ed. Garland, D.J. & Endsley, M.R., November 1995

Wiener, E.L. & Curry, R.E. "Flight Deck Automation: Promises and Problems", *Ergonomics*, Vol. 23, No. 10, p. 995-1011, 1980

Wong, G.A. "Development of Precision Runway Monitor Systems for Increasing Capacity of Parallel Runway Operations", *AGARD Meeting on 'Machine Intelligence in Air Traffic Management'*, May 1993

Yang, L.C. & Hansman, R.J. "Application of the Analytic Hierarchy Process for Making Subjective Comparisons Between Multiple Automation/Display Options", *Proceedings of the 6th IFAC/IFIP/IFORS/IEA Symposium on Analysis, Design and Evaluation of Man-Machine Systems*, Cambridge MA, June 1995

Appendix A Control Signal Flow Model of the Alerting System

Chapter 1 identified four components of an alerting task: Information Processing, Hazard Assessment and Alerting, Decision Making and Selection of a Resolution to the Hazard, and Control Actuation. Alerting systems are given responsibility for some (or all) of these components. This compartmentalization breaks down the alerting task into easily described concepts. Each component may require highly complex computations, but their conceptual separation allows each to be examined individually.

Using the control signal flow in an estimator-controller as a basis, a model of the combined alerting system and pilot can be made. This model is shown in Figure A.1 with the nomenclature used in control and estimation theory. As an open- or closed-loop controller, the combined alerting system and pilot drive the command inputs to the aircraft in order to reduce or eliminate a hazard and meet some performance constraints or specifications, measured at the aircraft's states.

This appendix will discuss the four components of the alerting task. Then, this representation of the alerting system as a dynamic controller will be used as a structure for describing the requirements on each component for an final satisfactory resolution of the hazard, and for evaluating the sensitivity of the alerting system to perturbations such as pilot non-conformance.

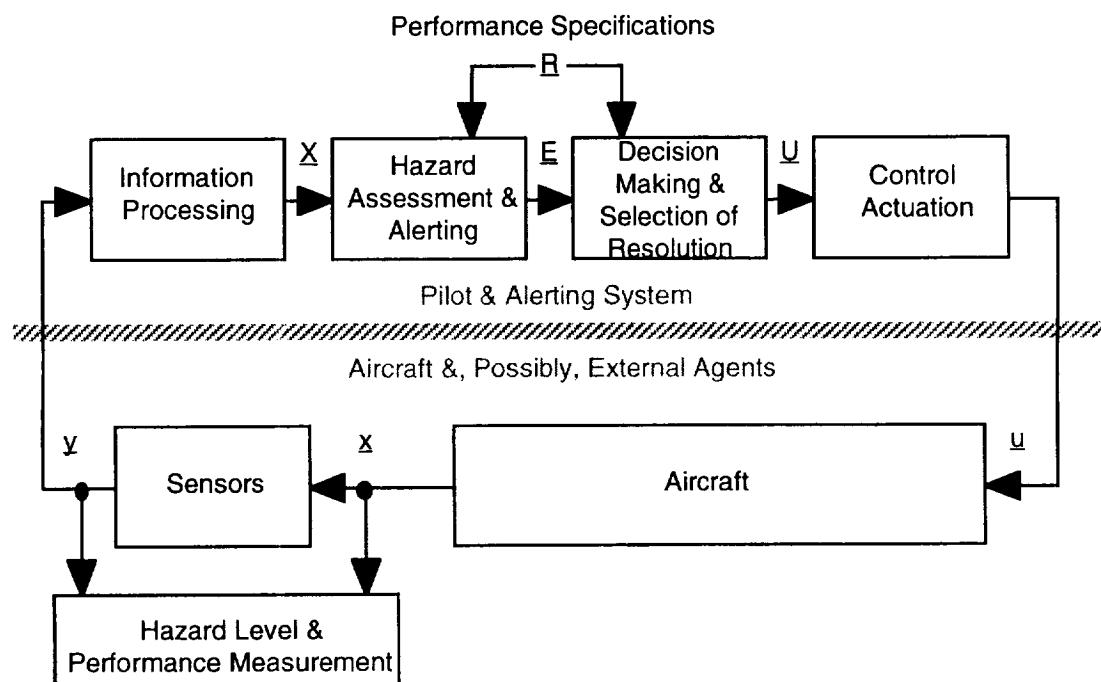


Figure A.1 Model of Controller for Time-Critical, Decision-Making Task

Appendix A.1 Variable Sets

Six variable sets describe the general control signal flow shown in Figure A.1. Two are given by physical properties of the dynamic system controls and of the sensors. The other three internal sets, on the other hand, are driven by the sixth variable set, the performance specifications.

Performance Specifications R The performance specifications list the required system attribute(s). These attributes include allowable hazard levels, and also the aircraft states allowed in the process of resolving the hazard.

The attributes may be expressed as extremal values or absolute limits on the hazard ('Do not get within 500 feet of the other aircraft'), relative values ('Climb faster than the other aircraft'), target values ('Use a bank of 30 degrees in the avoidance maneuver'), and objective functions for optimization ('Generate the greatest separation between yourself and the intruding aircraft'). These expressions can therefore provide constraints on the final solution or optimization criteria.

A difference should be noted between the overall system performance desired, aggregate over many iterations of the task, and the performance specifications the controller considers during the completion of a single incident of the task. For example, an aggregate measure of a collision avoidance system is the relative number of false alarms compared to the number of total alerts issued. During a single potential conflict, however, the controller's performance specifications may not necessarily directly consider the aggregate measure, but may instead provide fixed limits on the potential hazard acceptable in any single case. Identification of the performance specifications for each iteration of the control loop must be made during the automation design, human operator training or procedures development.

Projected Hazard E The projected hazard may be assessed directly from the current state estimates, or it may require projecting the system state estimates into the future some amount of time. The hazard may be represented by discrete values such as decisions to alert, or by continuous values.

Current State Estimates X The state estimates are the minimum set of current variables required to estimate the hazard. Depending on the system and the structure of the performance specifications, these variables may be the continuous dynamic variables found in traditional control analysis, or the variables may also include binary assessments and descriptions of discrete events. The states may include, in addition to current physical properties of the system, current knowledge of target information towards which the system is being steered by an independent inner-loop controller.

For example, alerting logic development has identified states required for good alerting decisions & avoidance maneuver selection, including knowledge about the intruder's position, altitude, heading, speed, bank, and vertical speed.

Sensor Outputs y The sensors from the dynamic system provide a set of sensor information. This sensor information may match exactly the form of the state estimates used by the controller, or may provide redundant or insufficient information. Associated with each of the sensor measurements are their properties including their update rate, latency, and statistical characteristics.

In the design of alerting systems, the available sensor information may drive the alerting system's functioning, or the desired automatic functioning and sensor requirements may be determined from the performance specifications. Established disciplines in instrumentation and estimation are available for analysis of the trade-offs involved in this analysis.

Dynamic System Control Inputs u The dynamic system control inputs are specified by the physical control actuators available, and the limits the operating environment and procedures place upon their use. The identification of the control inputs also depends on location of the controller; for example, an air traffic controller's inputs in an aircraft collision avoidance task are generally voice commands to pilots, whereas pilot's inputs in collision avoidance tasks use aileron and elevator commands.

Resolutions to the Hazard U In order to achieve the performance specifications R using the dynamic system controls available u , a set of possible hazard resolutions can be identified. The hazard resolutions are of a form that stipulates desired changes to the controller's internal states; therefore, they may or may not resemble the dynamic system control inputs. For example, collision avoidance maneuvers are often thought of in terms of desired pitch and/or vertical speed, and desired bank and target heading; these are separated from the aileron/elevator task performed by the control actuation.

Appendix A.2 Controller Sub-Functions in an Alerting Task

This section describes characteristics of the four, serial sub-functions required to generate the system control inputs u from the sensor outputs y in decision-making task, using collision avoidance as an example. Depending on the task and system characteristics, the sub-functions may require simple or complex operations, large amounts of controller memory, and detailed knowledge of the dynamic system characteristics and sensor accuracy.

Information Processing

In the case of collision avoidance, this task transforms the sensor information into the state estimates used for alerting and deciding on an avoidance maneuver, such as relative aircraft position and convergence rate. In the more general cases, this sub-function can be defined as the transformation of the sensor inputs y into the current state estimates X required for the controller to meet its performance specifications.

If the sensor inputs are of the same form of the state estimates and are sufficiently noise-free and accurate, then this transformation is simple. However, many characteristics of the sensor inputs may require more substantial computations, as shown in Figure A.2. Noisy or inaccurate

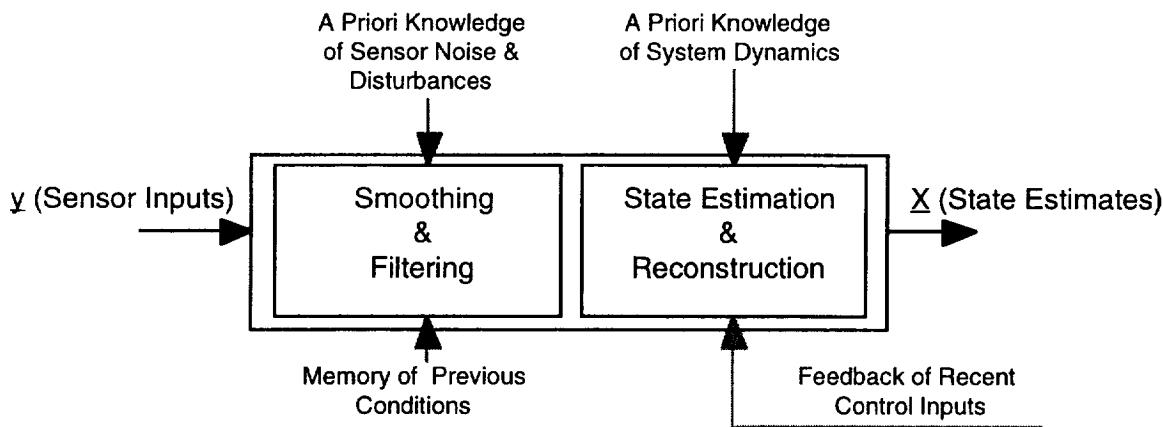


Figure A.2 The Information Processing Sub-Function

sensor inputs can require smoothing and filtering, which in turn requires before-hand knowledge of the sensor characteristics and the use of memory to keep track of previous conditions. Missing, inaccurate, or infrequently updated information, or information which provides the same measurements from more than one separate sensor, requires computations for reconstructing and estimating the states, which in turn requires before-hand knowledge of system dynamics and possibly knowledge of the latest control inputs made to the system.

Hazard Assessment & Alerting In the collision avoidance task this sub-function is responsible for making the binary decision to perform an avoidance maneuver. In a more general sense, this sub-function is defined as the transformation of the current state estimates X into estimates of the current or projected hazard E . As its name implies, this sub-function is comprised of two related evaluations.

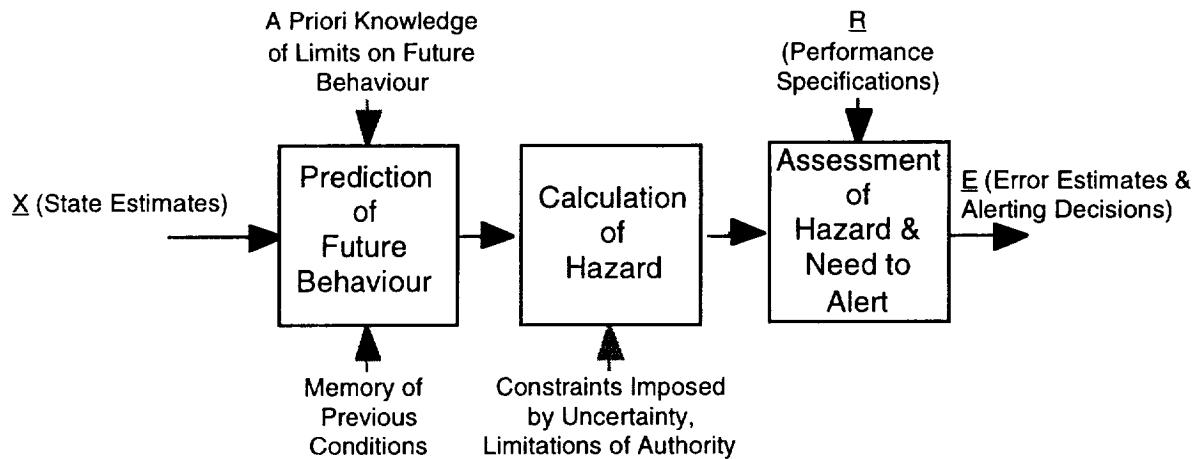


Figure A.3 The Hazard Assessment & Alerting Subfunction

1) Hazard Assessment. This component calculates the hazard levels from the system state. The calculation may be simple (when the states and performance specifications are of the same form) or it may require additional transformations of the data. For example, the determination of the closest point of approach between two aircraft requires computation of the inter-aircraft range from their current and predicted state (positions and velocities).

Prediction of the hazard levels in the future is required when the performance specifications require reactions to *predicted events*, before they actually occur. In the collision avoidance task, for example, this prediction is vital because of the need to identify the need for an avoidance maneuver with sufficient lead time to mitigate the effects of response delay and the time for the aircraft dynamics to be accelerated to the avoidance maneuver specifications. The time horizon of the prediction may be set by design, bounded by the performance specifications, or changed in real-time as warranted by the situation.

2) Alerting. Alerting can be seen as a binary switch which enacts the decision-making sub-task. For example, traffic alert systems generally provide alerts when the collision hazard passes a threshold; before the alert, the operator is not required to attend to the task, but after the alert, the operator is expected to focus on an immediate solution. This alerting function is generally a 'judgment call' based upon thresholds and/or goal assessments.

This sub-function is the simplest when the state estimates and, if necessary, its linear extrapolations can be directly compared to the performance specifications. However, many characteristics of the system may require more substantial computations, as shown in Figure A.3. Linear extrapolations of non-linear system dynamics may not have sufficient accuracy and would

then require greater before-hand knowledge of the dynamics. Unpredictable systems may require calculation of the envelope of possible actions. Incompletely measured system dynamics can require storage of past values. The assessment of hazard or ‘alerting logic’ may require probabilistic assessment, use of repeated sensor inputs to ascertain trends, and limits to the alerting horizon. The assessed hazard may be of a very different form than the states, and require substantial computation.

Decision Making. This sub-function, in the collision avoidance case, decides upon an avoidance maneuver which will meet the criteria given by the performance specifications. In the more general sense, it decides upon resolutions to the hazard \underline{U} , when required, to eliminate unacceptable hazards.

The simplest form of decision making has sufficient knowledge of the system dynamics, or the hazard has such obvious solutions, that a single decision can be quickly found through a simple transformation of the error estimates.

However, several factors can make the decision making sub-function more complex, as shown in Figure A.4. Evaluating the performance of the decisions involves knowledge and prediction of the system dynamics. Because the interrelationships between future hazard and hazard resolution are not always known with sufficient accuracy, an iterative search may be required; the number of iterations and the convergence of the search strategy depend strongly on the accuracy of the system

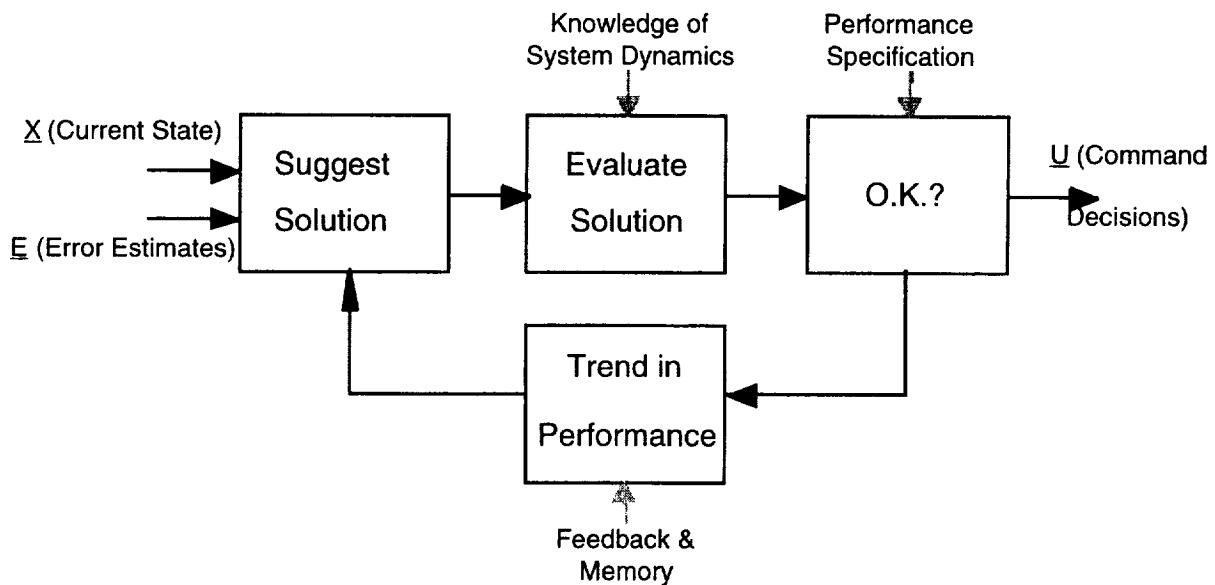


Figure A.4 The Decision Making and Selection of a Hazard Resolution Sub-Function

model and of the error estimates. The iteration may build upon known relationships; if it lacks such knowledge, computation of the predicted states may be needed as an intermediate step in the iteration or feedback may be required to identify the effects of decisions. Disturbances to the system can also require the use of feedback and updates to the commanded hazard resolution.

Control Actuation In the collision avoidance task, this sub-function describes the pilot or autopilot's task of controlling the aircraft through the avoidance maneuver, within the specific criteria, to the target states defining the avoidance maneuver. In the more general sense, this sub-function can be described as the generation of the system controls \underline{u} that will enact the commanded resolution to a hazard \underline{U} . The simplest form of control actuation is possible when the system controls are closely related to the commands and are therefore easy to compute. Detailed knowledge of the variables' relationship can also enable more streamlined calculations.

Like the decision making process, however, several factors may require an iterative search for the system controls. These factors include the lack of complete knowledge about the relationship between the variable sets, models which require a numerical solution search, and control strategies which rely on the use of feedback to evaluate and update the system controls.

Appendix A.3 Defining the Intermediate Requirements for Each Sub-Function of the Alerting Task

As alerting systems become more complex, the requirements for each sub-function (and its corresponding component in the alerting system) become more difficult to ascertain. For example, in the collision avoidance task, the ultimate performance specification is a requirement for safe aircraft separation at all times. The alerting system's performance at this metric may be sensitive to inaccuracy or delay at any junction throughout the alerting task. However, these sensitivities can be hard to determine and may be correlated.

The break-down of the alerting task into four serial sub-functions allows the alerting system to be modelled as a serial set of constraints, requirements and trade-offs, as shown in Figure A.5. In some cases, absolute constraints and requirements may be imposed on any of the alerting system's components due to the sensor inaccuracies, the dynamics of the aircraft, and the control inputs available. Beyond these absolute limits, trade-offs may also be found between sub-functions.

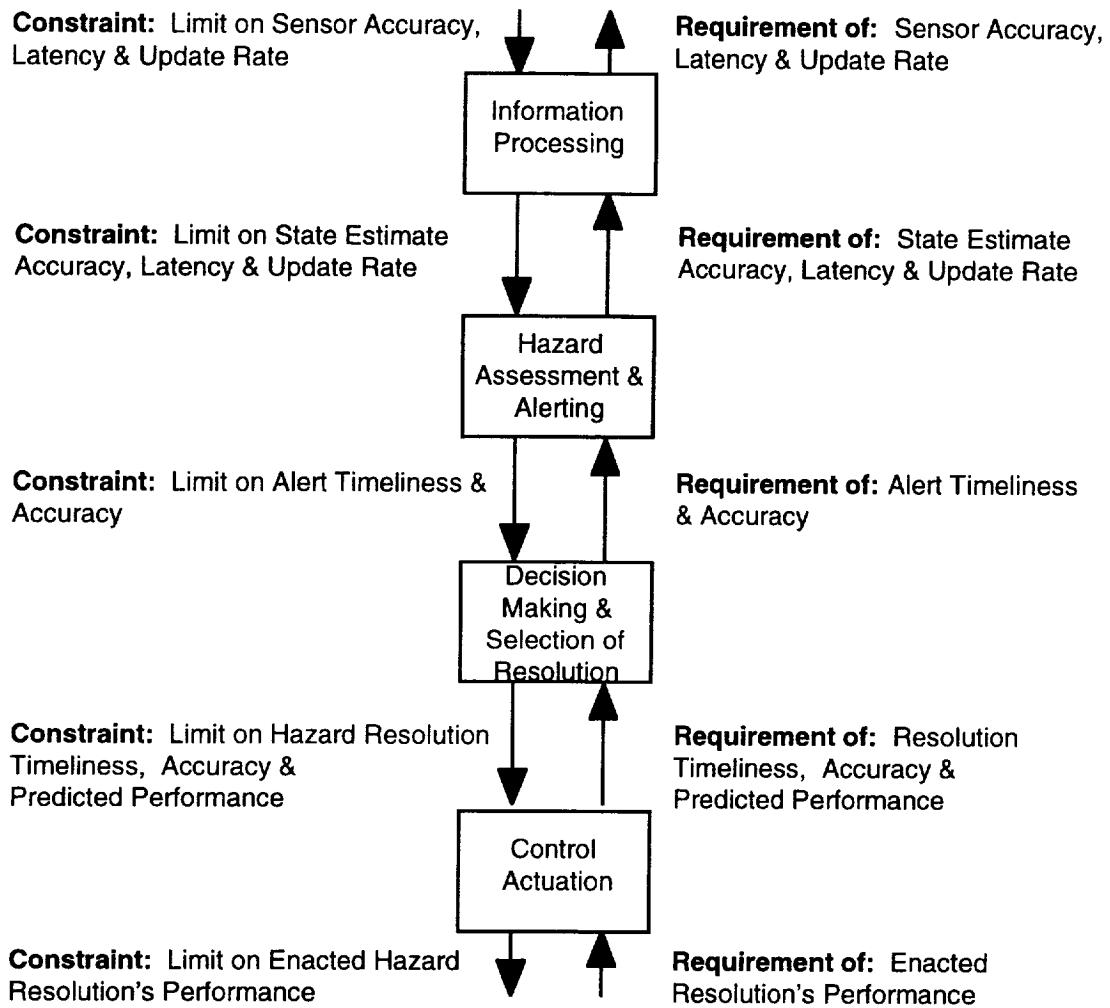


Figure A.5 Schematic of the Constraints and Requirements On Each Alerting Task Sub-Function

In the collision avoidance task, for example, the performance specifications may limit the loads placed on the aircraft, constraining the allowable control outputs; the performance specifications may weight highly certain avoidance maneuvers, constraining the range of decisions which may be allowed; the performance specifications may require a safe aircraft separation, placing an artificial criteria on the lead time required of the alerts; finally, to achieve this lead time, a certain accuracy may be required in the state estimates provided by the information processing.

This emphasis on the performance specifications has several implications for the model of the alerting task. They affect the workings of the alerting task components and the variable sets that must be generated. As such, even small changes in the performance specifications may result in

large changes in the structure and complexity of the components. Finally, the number of variables, and the precision with which they are calculated, depends on the required resolution of the performance specifications; crude performance can be achieved by simple models with small variable sets and reduced numerical precision, while exact performance may require inclusion of more variables and require more numerical precision.

This conceptualization of the alerting task may be beneficial in determining the effects of time-delay and variance in the pilot's execution of the alerting system's commands. It may also provide a structure for analyzing and designing alerting systems in general.

Appendix B Models of the Human Behavior at Alerting Tasks

Alerting tasks have both decision-making and aircraft-control components. This appendix summarizes the current literature on these two different areas of study.

Appendix B.1 Human Control of Dynamic Systems in Control Tasks

Many studies have examined flight handling qualities and the performance of humans as controllers of dynamic systems during continuous compensatory (or pursuit) control tasks. These types of control tasks differ from alerting tasks because they attempt to zero continuous error signals using rates and accelerations. Therefore, unlike alerting tasks, analysis of the system can use frequency-response methods to evaluate controller performance at a few well-defined objectives. Despite these differences between the tasks, examination of human performance at simple control tasks provide two valuable insights into human operator behavior during other tasks with decision-making components.

First, several studies have attempted to model exactly the input-output properties of the human controller at simple control tasks. Some properties of the human control capability have been well identified, and can be used as absolute limitations on human response time. For example, the Human Structural Model in Figure B.1 identifies the characteristics of signal flow through human vision (on the far left) and the dynamic properties of the human neuro-muscular system.

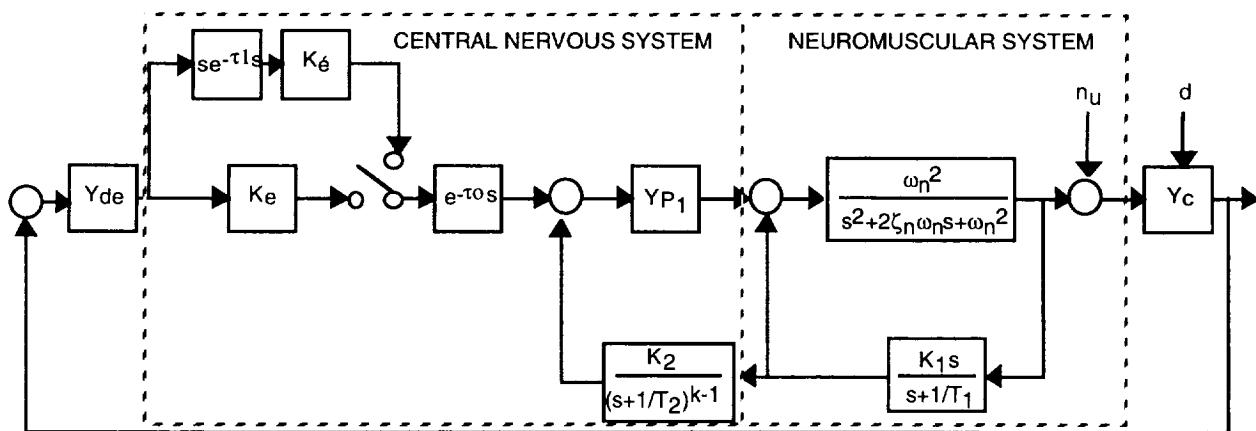


Figure B.1 Human Structural Model for Compensatory Control Tasks
(Redrawn from Hess, 1980)

The second insight from these models stems from the difficulty experienced in exactly specifying the components of the Human Structural Model which relate to the cognitive element of the signal processing; these elements have been found to vary depending on the task and the system being controlled. McRuer (1989) identified the tendency of the human controller to change properties so as to keep the dynamics of the entire system, *including both the aircraft and the human*, as a simple integrator with a time-lag. This variation by the human makes his or her own properties appear erratic but keeps the total system well-behaved. This model holds for the range of dynamic systems which humans can control, and starts to break down when the system response is too fast, is too unstable, or has too much time delay.

This ability of human controllers provides the largest single difficulty in analyzing the human controller at simple tasks with mathematical certainty. In alerting tasks, an analogous behavior is the ability of the human to switch strategies in an attempt to improve total system performance and keep the task manageable. Alerting tasks may, like simple control tasks, be easier understood by examining the constraints and influences on the human operator, rather than attempting to exactly predict their characteristics.

Appendix B.2 Complete Models of Human-Automation Controllers at Decision-Making Tasks

No proven method of completely describing and predicting the behavior of a human operator in a decision-making task has yet been surmised. However, a great deal of insight can be gained by examining the literature on decision-making to understand human behavior in alerting tasks.

In drafting a prescriptive model of human behavior and human interaction with automation, Johannsen (1992) identified several problems, including the problems involved with handling symbolic information (such as ambiguous or bounded goals) and knowledge-based information, and the ability of the human operator to change their own control strategies, making definition of the exact control paths and control structures difficult.

Other studies have instead used a similar breakdown of the task into sub-functions as used by this thesis' model. Some examples are listed in Table B.1. Although the different models have differing levels of detail, the emphasis is on understanding the system state, monitoring the system state or assessing hazard levels, evaluating the need for action, making decisions, and execution the control decisions. Cacciabue (1993) additionally includes memory as a higher level 'meta-process'.

Table B.1 Examples of Task Breakdowns Used in Other Studies

Source	Sub-Functions
Knaeuper & Rouse, 1985	1) Recognition and Classification (of the problem) 2) Planning (of a solution to the problem) 3) Execution (of the solution) and Monitoring
Pawlak, Brinton, Crouch & Lancaster, 1996	1) Planning 2) Implementation 3) Monitoring 4) Evaluation
Cacciabue, 1993	1) Observation/Perception 2) Memory (recall of information) 3) Interpretation (identification/diagnosis) 4) Choice (planning / decision making) 5) Execution of a plan.
Rasmussen, 1986	1) Detect the need for intervention 2) Observe some important data 3) Identify the present state of affairs 4) Evaluate their possible consequences 5) Define task 6) Formulate procedure to achieve task 7) Execute task.

Some variations do occur between the models. These may be influenced by the type of systems and types of tasks for which they were developed. For example, Rasmussen (1986) looked at the diagnostic task of identifying the cause of an anomaly *once its presence was known*. This may explain the differing order of the sub-tasks, the detail to which they are broken down, and the emphasis given to each.

Studies have also emphasized the importance of the human operator's mental model of the system (For example. Rasmussen, 1986). The term 'mental model' has been found to relate to several different components of the control signal and controller. A mental model of the situation may include the dynamic variables (defined in Appendix A as the Current State Estimates, Estimated Hazard, and Hazard Resolutions) which evolve throughout the task. The mental model may also include the *a priori* knowledge or representation of the system used by the human operator in evaluating the dynamic variables.

Appendix B.3 Studies of Human Execution of the Alerting Task Sub-Functions

Many studies have examined specific sub-tasks of the control process. Although their results are often application specific and are usually expressed in qualitative or conceptual terms, many similarities with the components of the alerting task used by this thesis can be found.

The concept of *Mental Model-Interface Compatibility* emphasizes the importance of identifying the internal state representation used by the human operator during the task, and of presenting the correct information to the operator to maintain this mental model. (Smolensky, 1995) This concept has obvious analogies to the identification of the Current State Estimates and the sub-function 'Information Processing'.

The concept of situation awareness has also been widely studied. Endsley (1995) has defined three levels of situation awareness (SA): Level 1 SA, the perception of elements in the environment; Level 2 SA, the comprehension of the current situation; and Level 3 SA, the projection of future status.

The sensor outputs identified by this thesis' model correlate with Level 1 SA, as both refer to the environment variables which can be perceived by the human operator. These two concepts differ, however, in that this thesis identifies the sensor outputs that are physically available, while Level 1 SA refers to the sensor outputs that are perceived by the human; for a variety of reasons these two sets may not be identical.

Level 2 SA is defined to go "beyond simply being aware of the elements that are present to include an understanding of the significance of those elements in light of pertinent operator goals." (Endsley, 1995) This definition suggests Level 2 SA is comprised of the values calculated in the Information Processing sub-function for use in the Hazard Assessment; in the collision avoidance task, Level 2 SA would include knowledge of the other aircraft's position relative to the own aircraft.

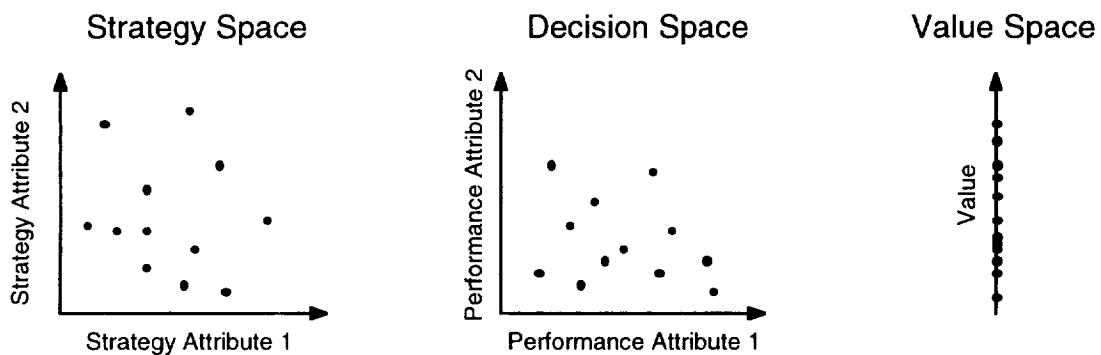
Finally, Level 3 SA is defined as "the ability to project the future actions of the elements in the environment... [SA] includes comprehending the meaning of that information in an integrated form, comparing it with operator goals, and providing projected future states of the environment that are valuable for decision making." (Endsley, 1995) These values are analogous with the projections of the hazard and need for an alert. For example, in the collision avoidance task, Level 3 SA would include the predicted collision hazard.

Some debate has occurred over whether SA is a product or a process. (For example, Billings, 1995; Sarter & Woods, 1995; Wickens, 1995) Accepting the definition of SA a product -- the

variables passed between sub-functions in the alerting task, for example -- the process of generating this knowledge, sometimes referred to as 'Situation Assessment', is represented by this thesis' model as the controller's sub-functions Information Processing, and Hazard Assessment and Alerting.

A 'disconnect' has been noted between SA and decision making. Although many poor decisions have been attributed to poor SA, Endsley (1995) notes "Individuals can still make poor decisions with good situation awareness". This effect may be related to the serial nature of the components of the alerting task; although a good estimate of hazard and system state is required as its inputs, the decision-making sub-function may be too difficult for successful resolution.

Other studies have examined the decision-making process more directly. Patrick (1996) described the decision process as having the three states shown schematically in Figure B.2. First, the decision maker generates a set of strategies (prototype command decisions) which can be located in 'strategy space' by their properties. These strategies are each evaluated for their value in each performance attribute, which identifies each in 'decision space'. If the relative weighting of these performance attributes can be determined, then the decision space can be collapsed to a one dimensional 'value space'; Patrick used utility theory to model this final mapping. The strategy with the highest value would then be picked as the 'best'.



**Figure B.2 Patrick's Model of the Decision Process
(Attributes of Possible Decisions are Plotted by Dots in Each of the Three Different Attribute Spaces)**

The concept of Cognitive Task Analysis seeks to identify the strategies available to, and used by, the operator. Several means are used for cognitive task analysis, including "interviews, cognitive probes of critical incidents, presentation of challenging simulations, ... ". (e.g. Thorsden

et al) As such, this method is often used for testing developed systems, rather than as a predictive assessment of possible future systems.

Other methods have modelled human decision-making by generating equivalent expert systems or rule-based structures which operate on logical rules of inference. These models need to be extensive, however, to account for the variations in human strategy, and presume consistent decision paths. (e.g. Knaueper & Rouse, 1985)

Appendix C Descriptions of Alerting System Roles

Chapter 1 briefly described two criteria for defining the role of an alerting system. First, the sub-functions which the alerting system performs can be identified. Second, the supportive or executive role of the alerting system could be identified by examining whether it is intended to provide the pilot with more information to perform each sub-function themselves, or to negate the need for the pilot to perform the sub-function.

This appendix describes these alerting system roles in greater detail. Then, the limitations of these definitions are discussed in section C.6.

Appendix C.1 Fully Manual Control

Fully manual control, as shown in Figure C.1, does not use the alerting system for any assistance at any function. The pilot is given only the raw sensor outputs, and is responsible for all of the sub-tasks required to generate the control inputs to the dynamic system. This role of alerting system is representative of traditional control schemes for continuous control tasks.

Appendix C.2 Automatic Information Processing

Automatic information processing, as shown in Figure C.2, involves the automatic system in the information processing sub-task. The alerting system then provides a higher-level display of the current system state to the pilot, and the pilot performs the remaining sub-functions without assistance.

Conceptually, it can be difficult to distinguish between information processing by an alerting system and the inclusion of better or more appropriate sensors. Presentation of the current information required by the pilot can be accomplished by correct sensor implementation, i.e. by designing the system to include a sensor suite whose outputs exactly match the pilot's required current state estimate; from the pilot's point of view, such a system might be indistinguishable from one with poor sensors and a great deal of automatic information processing.

In addition, with highly complex sensor information, the distinction between sensor functions and alerting system functions can become blurred; for example, in the task of processing and presenting radar returns of aircraft position, the decision to delineate between sensor functions and alerting system functions can fall anywhere between the antennae and the traffic display.

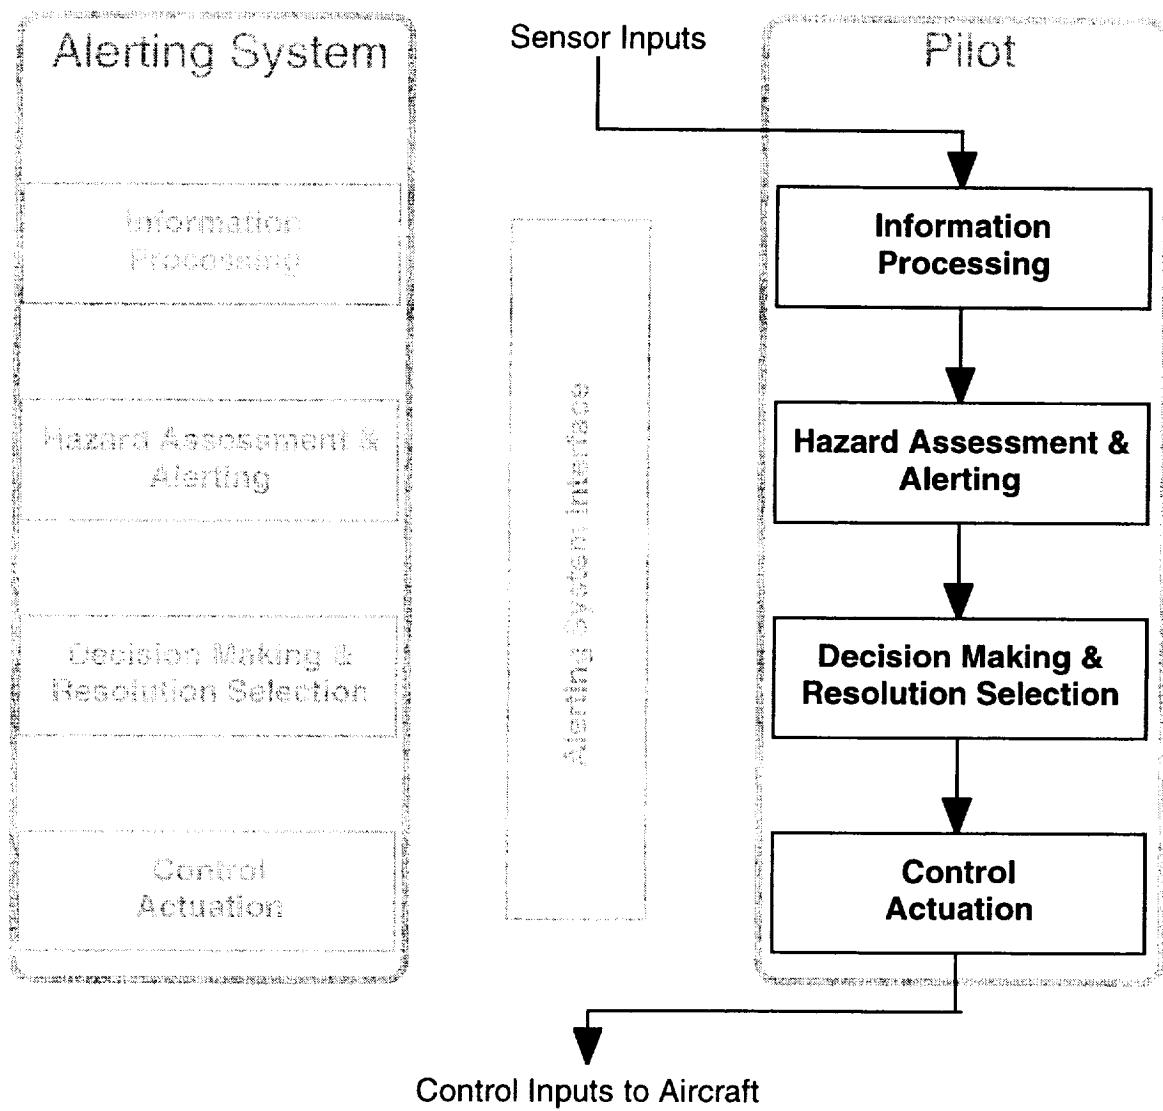


Figure C.1 Fully Manual Control

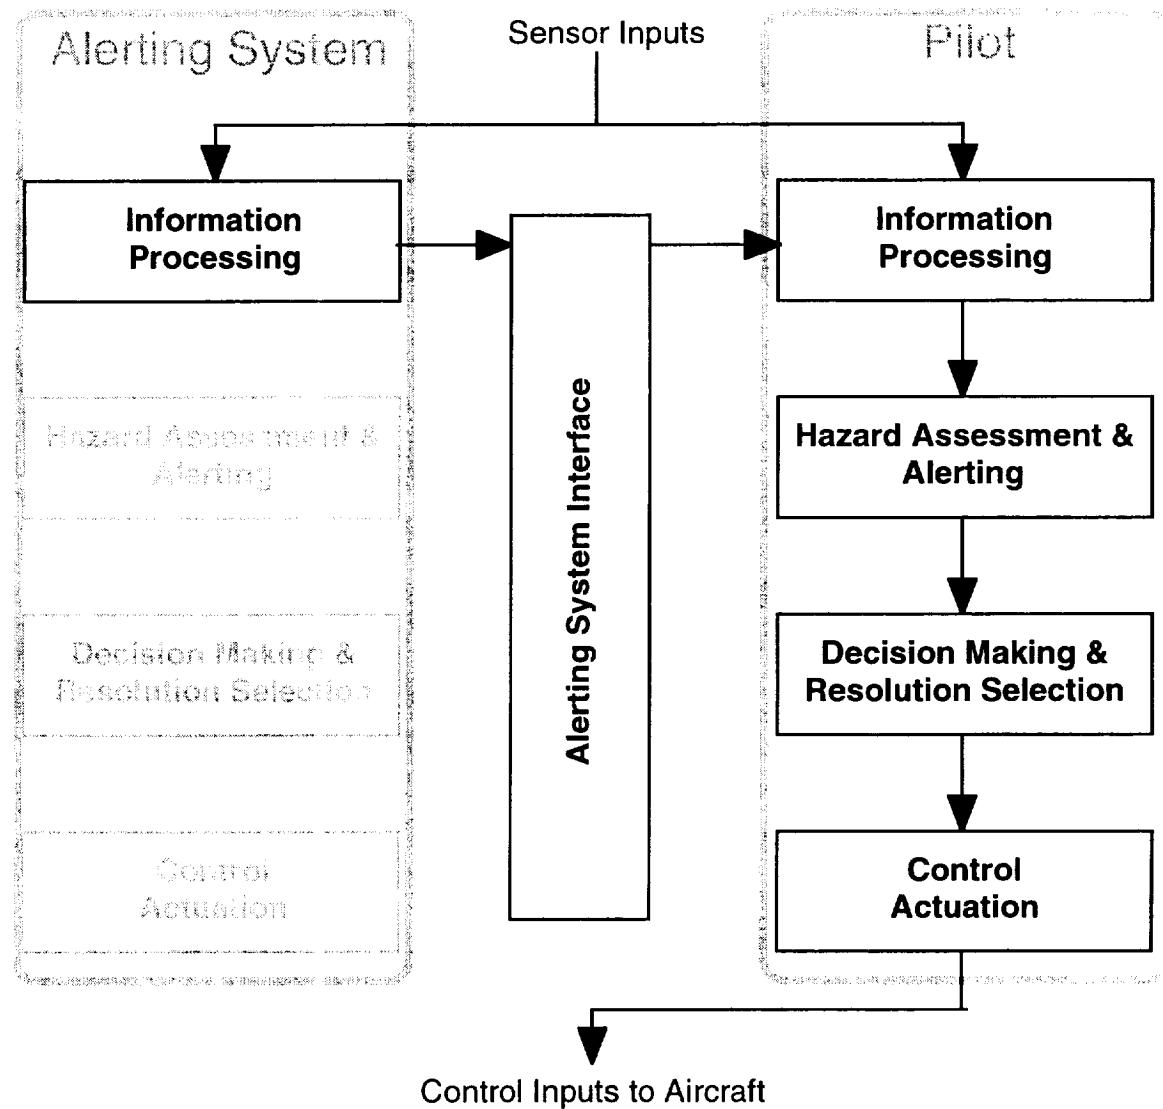


Figure C.2 Automatic Information Processing

In practice, these demarcations are often provided by the capabilities of established hardware components. The designer of alerting system can then think of automatic information processing as any processing done to data available from physical sensor units or system-wide data buses.

This level of alerting system involves two related features. First, the algorithms and hardware for automatic information processing are required. Equally important is the display to the pilot of the processed information in a way which makes the pilot's information gathering task easier than that experienced if the pilot was responsible for the information processing his or herself. The display of processed information, which may be multi-dimensional and involve a large amount of data, is currently the topic of many studies, both for alerting systems and at a conceptual level.

Appendix C.3 Automatic Hazard Assessment & Alerting

Automatic Hazard Assessment and Alerting, as shown in Figure C.3, additionally gives an alerting system partial or total responsibility for the two components of the hazard assessment and alerting sub-task. A 'Supportive' system can display to the pilot projected hazard information or future conditions, while an 'Executive' system generates executive alerts which the pilot must follow, with or without a display of the supporting information.

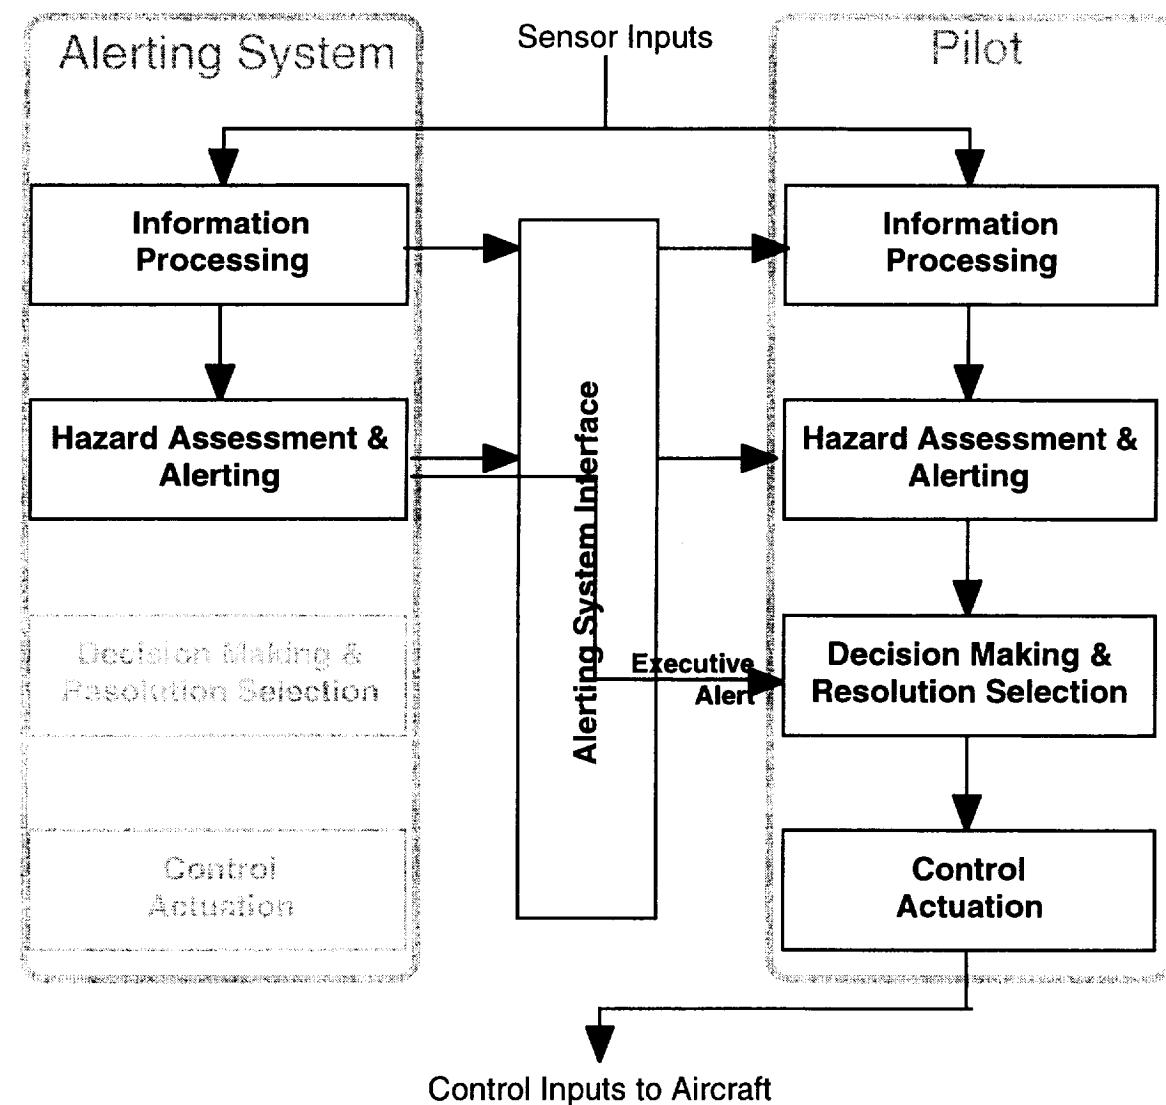


Figure C.3 Automatic Hazard Assessment and Alerting

As with automatic information processing, this role of an alerting system requires both an underlying computational component which generates the variables of interest, and an interface component which displays these variables. In the case of collision avoidance for closely-spaced parallel approaches, alerting logic is being prototyped for evaluation -- i.e. the underlying computational features are being developed. The alert displays, to date, have tended to provide simple indications of alerts without the display of predictive or hazard information, with the intention of relieving the pilot of the need to compute the need for an alert.

Appendix C.4 Automatic Decision Making & Selection of a Resolution to the Hazard

Automatic Decision Making & Selection of a Resolution to the Hazard, as shown in Figure C.4, gives partial or total responsibility for decision making to the alerting system. As noted in Appendix A, the decision making sub-task can vary greatly in difficulty, depending on how well the relationship between the potential hazards and the command decisions is known. Therefore, the alerting system can assist in the decision-making task at many levels; discrete examples in what may be seen as a continuum include:

- *Supportive Role:* A simple form of decision aiding can be provided by the automatically displaying the trends in the relationship between projected hazards and recent command decisions. This provides the pilot with the ability to interpolate the effect and required magnitude of further command decisions, without requiring the alerting system to have any knowledge of the system. For example, an alerting system may help a pilot select a turning avoidance maneuver by indicating the derivative of projected miss distance with respect to heading. This technique can make the pilot's component of the decision simple, but requires good knowledge of the system and is difficult to communicate to the pilot when the command decisions and/or the performance specifications are multi-dimensional.
- *Supportive Role:* The alerting system can compute and display the projected system states and hazards resulting from a selected command decision before it is activated and/or before it takes effect. When the effects of control decisions are difficult for a pilot to compute, this provides a quick decision aid which leaves the authority with the human. However, this technique requires the alerting system to have good knowledge of the system dynamics.

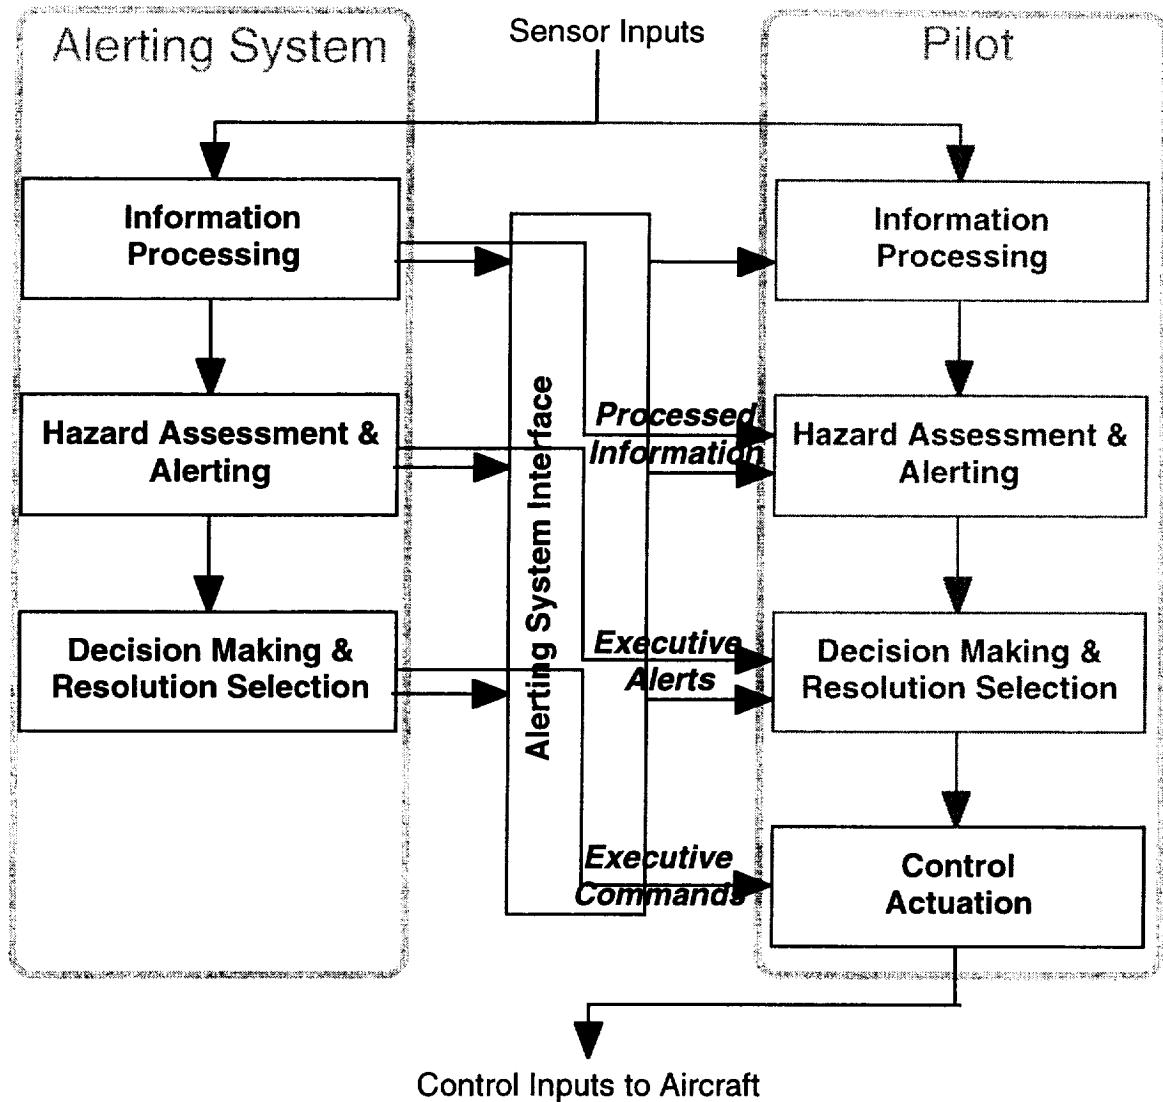


Figure C.4 Automatic Decision Making

• **Supportive Role:** The alerting system can present several decision options for the pilot to select from, with some illustration of their respective performance characteristics. This technique is suitable for tasks with multi-attribute, perhaps ambiguous, performance specifications for which the alerting system can not select a 'best' alternative. However, the alerting system requires the ability to not only evaluate command decisions, but to search for them.

• **Executive Role:** A fully automatic decision making system presents a command decision for the pilot to execute. The command decision may be in a conceptual form or may be very specific about the control actuations it requires. Current collision avoidance systems tend to be of this form, presenting flight director or other command state displays to the pilot. Supporting information may or may not be shown.

Appendix C.5 Fully Automatic Control

Fully automatic control, as shown in Figure C.5, gives the capability for all sub-tasks, including control actuation, to the alerting system. The pilot may or may not have access to complete information about the system states and automatic processes, as well as the ability to intervene or disconnect the alerting system from the control inputs.

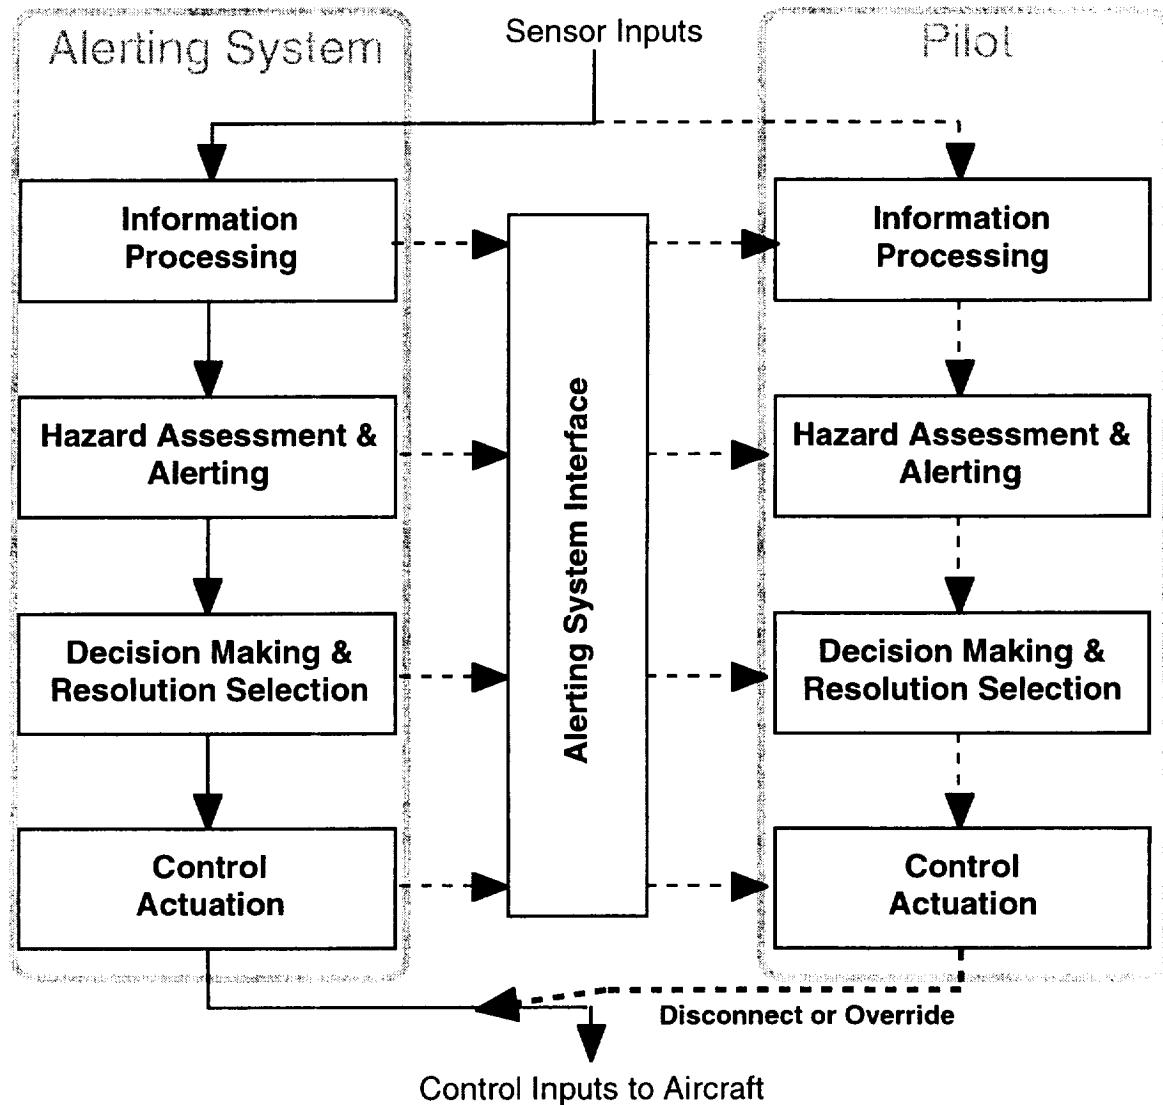


Figure C.5 Fully Automatic Control

Appendix C.6 Limitations of the Definitions of Alerting system Roles

These definitions of various roles for alerting system have limitations. For example, these definitions assumed an all-or-nothing use of the control paths. In reality, the alerting system may operate on some segment of the total control signals required to complete the task, with the remainder being processed by the pilot. To define such a case would require listing the role of alerting system for each of the possible control signals which, while possible, may result in an attention to details which loses sight of the descriptive power of these definitions.

In addition, these definitions of the roles of alerting system gave the alerting system responsibility for the earlier components of the alerting task; i.e. the control signals could flow down from the alerting system to the pilot during the process, but never up from the human to the alerting system. This generalization was made for several reasons.

First, in time-critical situations, such as collision avoidance, requiring the pilot to perform the initial sub-tasks such as information processing and hazard assessment and alerting, and to then enter the data through an interface to the alerting system for automatic decision making is often not possible within the time-available.

Second, such a relative breakdown of responsibilities is hard to distinguish from fully manual control of a system with modified control actuators. In the same way that information processing can be hard to distinguish from sensor functions (discussed in Section 2), programming an automatic decision-maker can be hard to distinguish from controlling a plant with more sophisticated control inputs.

Third, for most tasks, in order for the alerting system to perform the later sub-functions such as decision making and control actuation, it generally must also first perform the earlier sub-functions such as information processing and projecting, monitoring and alerting. Once the cost of developing the alerting system to perform these functions has been invested, it may be convenient to have the alerting system perform all the functions it is capable of, as a fully (or nearly so) automatic controller.

Finally, humans are generally considered better at judgment tasks than alerting system, while the alerting system is better at extensive calculations and monitoring. Therefore, alerting system development to date has tended to use alerting system as a supporting system to human decision making.

Of course, examples can be found where necessity dictates a switch in the roles of alerting system and pilot. One such example is a situation where automatic information processing may not be able to perform as well as human visual pattern recognition; a role of alerting system for this

case can be envisioned where the pilot identifies a target, or an anomalous feature of the target, to the alerting system and it makes decisions based upon the specified characteristics.

The definitions of the roles of alerting system also presumed the pilot continues to have available both sensor information and variables sets describing the alerting system's functioning. While this presumption does not necessarily need to be true for all cases, the availability of this information is generally considered valuable: it allows the human to remain involved in the task, helps with user acceptance of the alerting system, and allows the pilot to monitor, and possibly disconnect, the alerting system.

Appendix D Description of Traffic Alert and Collision Avoidance System (TCAS II)

Appendix D.1 Overview

Collision alerting and avoidance systems for air transport aircraft have been proposed and tested for several decades. (e.g. Andrews, Senne & Koegler, 1978) The Traffic alert and Collision Avoidance System (TCAS II) was mandated for installation in all passenger carrying aircraft with more than 30 seats by December 31, 1991. This system's role has been defined "as a supplement to the separation services provided by Air Traffic Control (ATC) and the 'see and avoid' concept." (Ciemer et al, 1993) As such, primary responsibility for aircraft separation remains with ATC. Pilots are expected to follow ATC instructions unless an executive Resolution Advisory (avoidance maneuver) is commanded by TCAS.

Recently, TCAS has been evaluated for two originally unintended uses -- as a means for overtaking other aircraft in trans-oceanic flight beyond the limit of Air Traffic Control radar surveillance, and as a means of ensuring aircraft separation during parallel approaches.

Appendix D.2 System Components

As shown in Figure D.1, the TCAS system includes several components. Sensors provide information about the neighboring aircraft to the pilot through a traffic display, and to the alert generation and avoidance maneuver generation logic. When an alert is issued and an avoidance

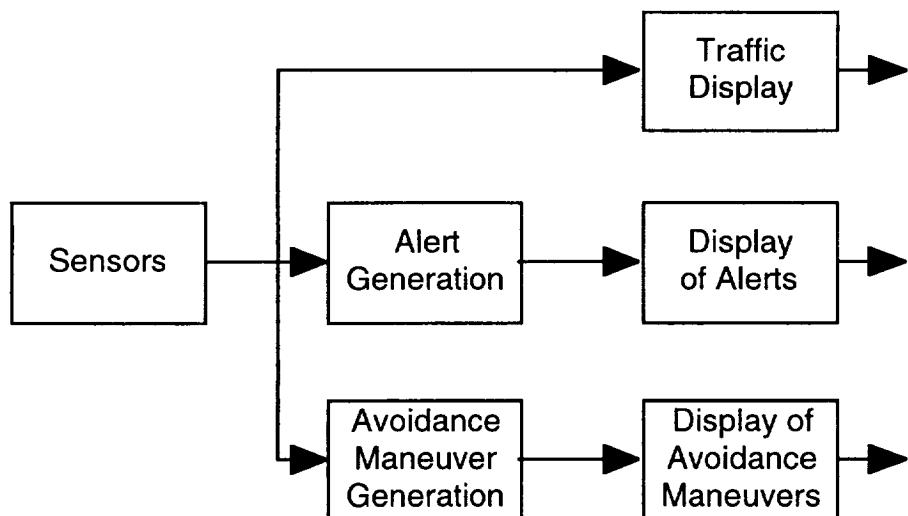


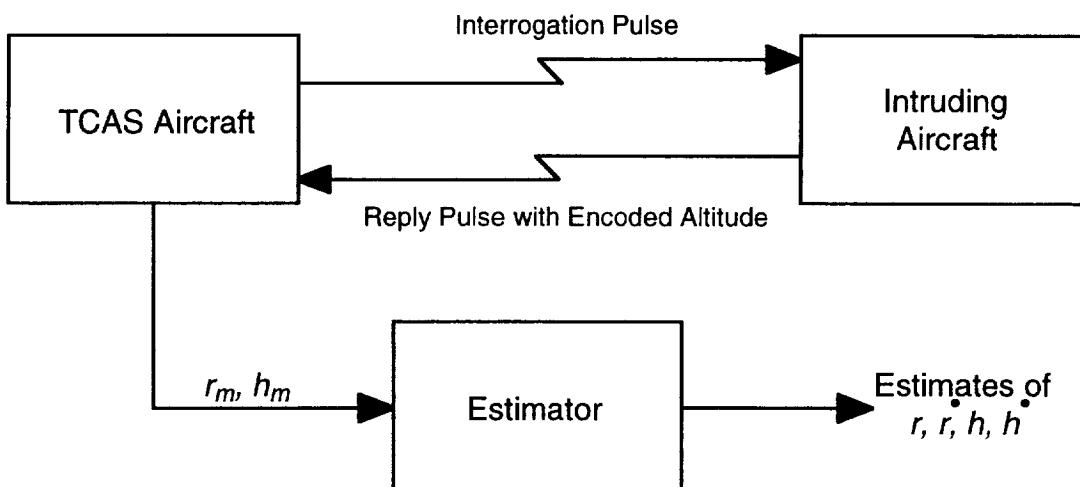
Figure D.1 Schematic of TCAS II Components

maneuver is commanded, the alert and avoidance maneuver commands are also displayed to the pilot.

Sensors

The TCAS sensors on each aircraft provide direct measurements of inter-aircraft range, r_m , and bearing to the other aircraft ψ . As shown in Figure D.2, these measures are made directly through interrogation radar pulses sent to, and returned by, neighboring aircraft. The radar pulse returned from the other aircraft also includes encoded information, giving the other aircraft's altitude, h_m (discretized to the nearest 100 feet) and, when an avoidance maneuver is in progress, whether the other aircraft is climbing or descending.

The range and altitude measures are then processed in an estimator. The range and altitude measures are filtered to reduce the effect of noise in the measurements. Range rate and altitude rate are inferred from the time history of range and altitude. The estimates of range, bearing and altitude are used for the display of the current traffic situation to the pilot. The estimates of range, range rate, relative altitude and altitude rate are used by the alert and avoidance generation logic.



**Figure D.2 TCAS II Sensor Information
(From Kuchar, 1994)**

Alert Generation

TCAS generates two levels of alerts -- a cautionary Traffic Advisory (TA) and an executive Resolution Advisory (RA). (RTCA, 1983) For an alert to be generated, two tests on the current state estimates must both be passed. Both are based on the predicted time to collision extrapolated from the current distance and rate estimates. The first test is calculated in the horizontal plane; the second test is calculated in the vertical plane. Figure D.3 shows a schematic for the horizontal test; the vertical test is similar. Both the TAs and RAs are generated by these types of tests. However, the two tests use different threshold values; TAs are generated with a larger projected time to collision, and are therefore generated before the RAs and commanded avoidance maneuvers.

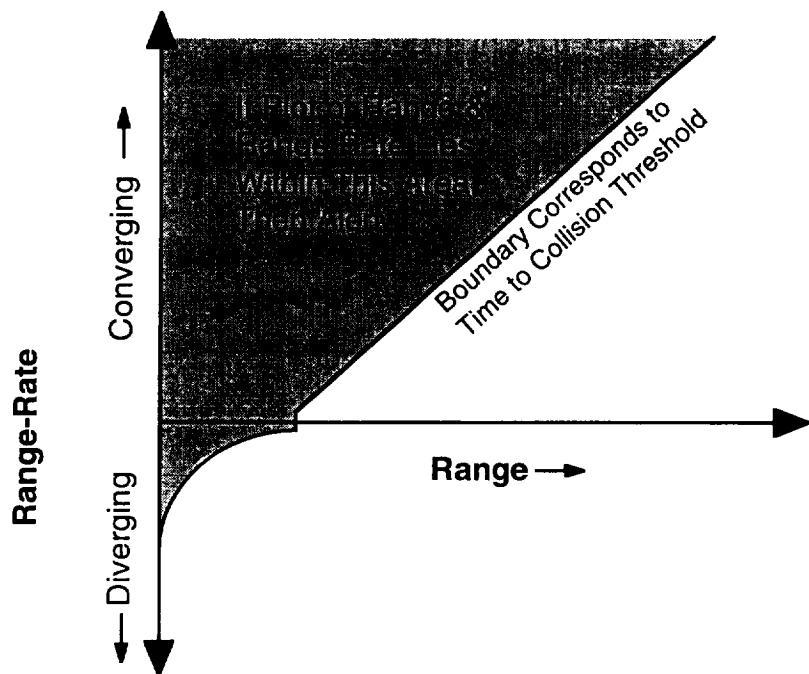


Figure D.3 TCAS Alerting Criteria, Based on Inter-Aircraft Range and Range-Rate

The alerting thresholds depend on aircraft altitude above ground and above sea level. An aircraft on a parallel approach may traverse two or three 'Sensitivity' levels and may be at a different sensitivity level than an aircraft it is in conflict with. Typically, a TA will be generated on final approach with 25-30 seconds remaining to a projected loss of separation; an RA will be generated with 15-20 seconds remaining. Additionally, TCAS alerts are automatically inhibited to only the cautionary Traffic Advisories (TAs) below 1000' above ground.

Avoidance Maneuver Selection

Once the alerting mechanism has generated a Resolution Advisory, an avoidance maneuver is selected. First, the two aircraft involved agree upon the maneuver 'sense' -- which aircraft will descend and which aircraft will climb. This negotiation is conducted through the encoded pulses used for range measurements. Once the sense of each aircraft's maneuver is decided upon, each TCAS unit, from this point, assumes the other aircraft will either maintain its current trajectory or execute an avoidance maneuver of the negotiated sense.

Each TCAS unit examines a specific set of avoidance maneuvers, as shown in Figure D.4. The 'weakest' maneuvers -- i.e. the maneuvers that require the least change to the own aircraft's trajectory -- are examined first. If a weak maneuver is predicted to achieve sufficient separation, then it is selected. Weak maneuvers include 'Preventive' commands to the pilot, such as 'Do Not Climb'. These maneuvers may not require the pilot to alter his or her trajectory, but instead limit any changes. If weaker maneuvers are not predicted to give adequate separation, then progressively stronger maneuvers are evaluated. In Figure D.4, for example, the weakest maneuver projected to create the minimum desired miss distance is the 'Don't Descend' command, so it is the maneuver displayed to the pilot. Maneuvers which require the pilot to change his or her trajectory are called 'Corrective', and a five second reaction time is expected from the pilot.

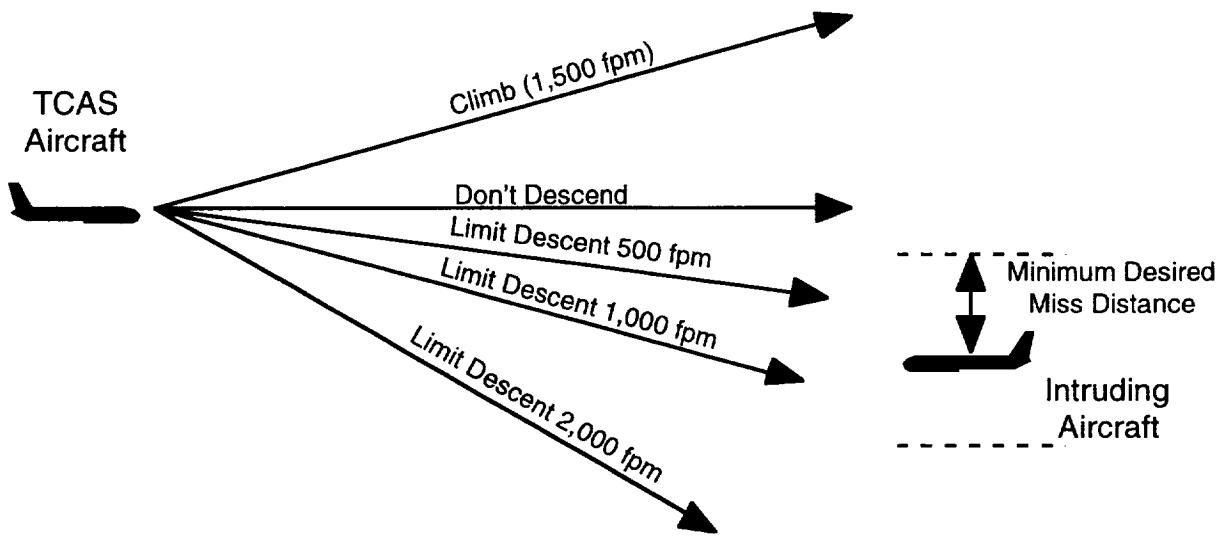


Figure D.4 Schematic of Maneuvers Evaluated (Climb Sense)

Pilot Interface and Reaction Times

The TCAS system displays three levels of information to the pilot -- the current traffic situation, alert information, and commanded avoidance maneuvers. The exact displays vary different airlines and in different aircraft. However, several features of the displays are standardized.

Figure D.5 illustrates a current implementation of the TCAS II traffic display, which is based on the moving map or Electronic Horizontal Situation Indicator (EHSI). This display presents navigation information to the pilot in the horizontal plane from a top-down view. The own aircraft position is indicated by the fixed triangle at the bottom center of the screen. The position of navigation features are drawn relative to the own aircraft symbol. For example, in Figure D.5 the landing runway, labeled '18R', is shown 10 miles in front of the own aircraft. As the pilot continues the approach, the runway will appear to move down the screen.

In all implementations of TCAS traffic displays, the other aircraft's position is drawn relative to the own aircraft symbol. The relative altitude, discretized to hundreds of feet, is shown by text. The relative position of the text, above or below the traffic symbol, provides an additional indicator of relative altitude. When the aircraft climbs or descends beyond a threshold rate, a vertical trend arrow appears next to the aircraft's symbol.

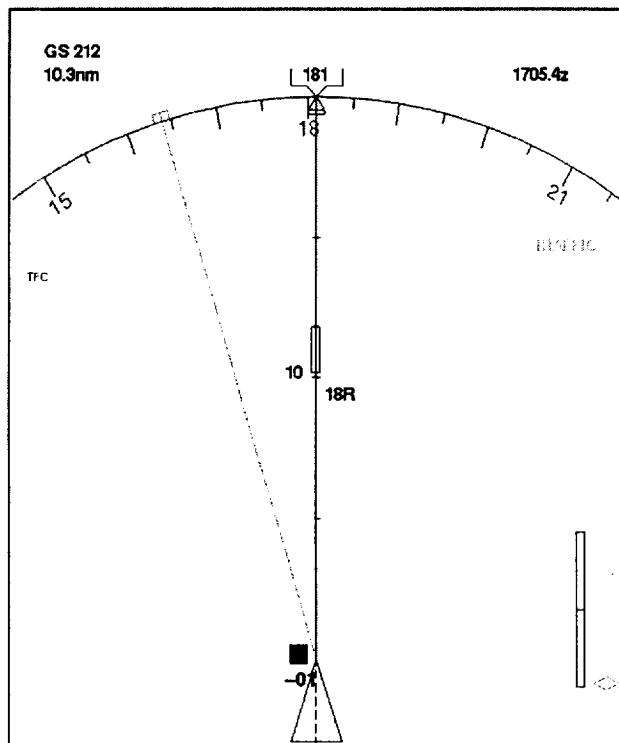


Figure D.5 TCAS II Cockpit Traffic Display, Integrated onto Navigation Display

The shape of the traffic symbol is also standardized across different implementations of TCAS. Normally a hollow white diamond, it becomes a filled white diamond when the traffic is within 6 miles horizontally and 3000 feet vertically of the own aircraft; during parallel approaches, the traffic almost always falls within this criteria. When the TCAS II system generates a cautionary Traffic Alert (TA), the symbol becomes a yellow circle. When the TCAS II system generates an executive Resolution Advisory (RA), the symbol becomes a red square.

In addition to the changes in symbology on the traffic display, the alert and avoidance maneuver information is presented by aural alerts, and by commands indicating the required vertical avoidance maneuver. Figure D.6 illustrates the presentation of the command information used by one implementation. When an RA is issued, a pitch command appears on the Primary Flight Display. To be in conformance with the command, the pilot must pitch the aircraft to keep the Aircraft Attitude Indicator outside of the area delineated by the Pitch Command. The avoidance maneuver is assumed to only be vertical; the pilots are expected to not also use a horizontal (turning) maneuver.

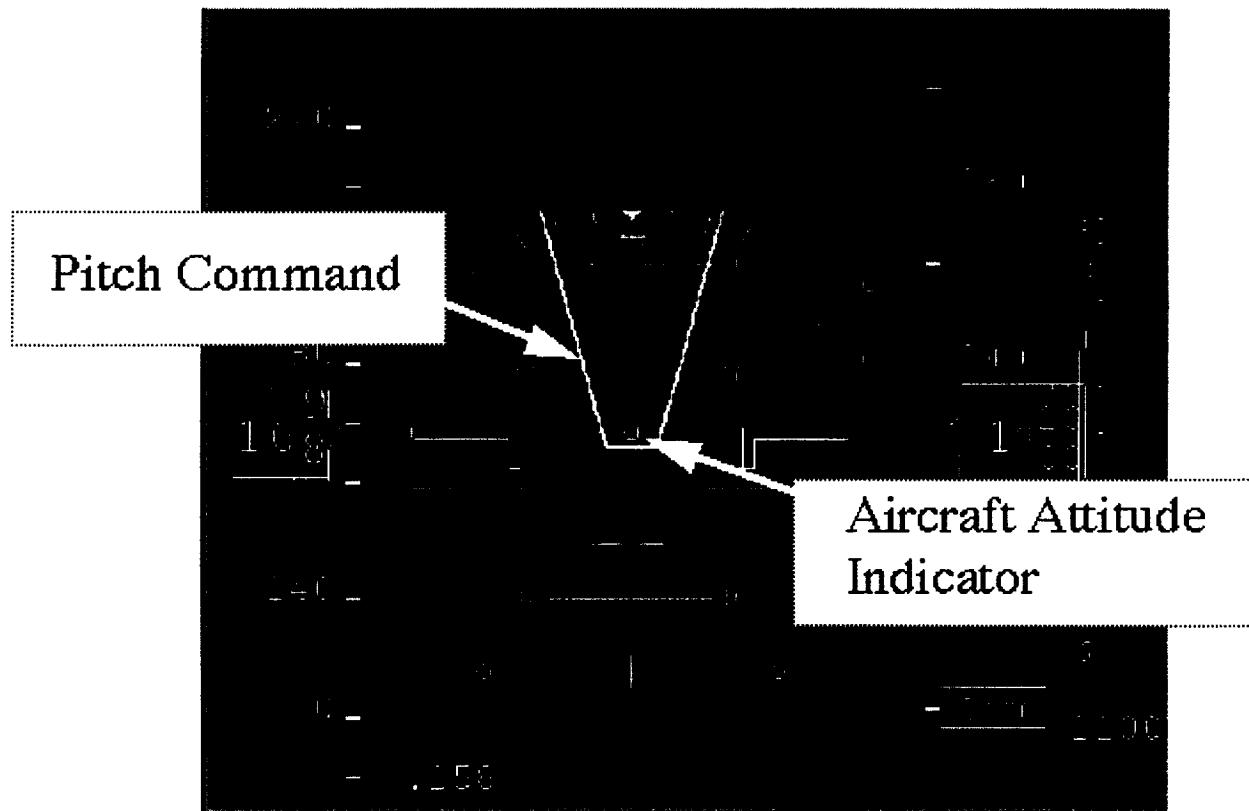


Figure D.6 TCAS Presentation of Avoidance Maneuver on Primary Flight Display

Appendix D.3 Predicted Performance of TCAS II During Closely Spaced Parallel Runway Operations

Several studies have evaluated the use of TCAS II for collision avoidance during closely spaced parallel approaches. Three issues appear to warrant further consideration.

First, although the alerts and avoidance maneuvers commanded by TCAS are predicted to generate safe aircraft separation, modifications to the TCAS logic may be required. Currently, TCAS Resolution Advisories and maneuvers are inhibited below 1000' above ground. This limitation reduces TCAS' ability to ensure separation for approximately the final three miles of the approach.

Second, numerical analysis predicts, with TCAS alert logic version 6.04A, a negligible number of false alarms at a runway spacing of 3000'. However, the number of false alarms increases with decreased runway separation. For example, with a runway separation of 1800 feet, the false alarm rate is predicted to reach 10%. (Folmar, Szebrat & Toma, 1994) Each false alarm requires two proximate aircraft to execute missed approaches. Frequent false alarms effectively reduce the benefit of allowing independent, simultaneous approaches. In addition, the rate of false alarms may degrade the immediacy of pilot trust in the system; for example, "Questionnaire data from pilots and observers indicate that low-altitude RAs continue to degrade pilot confidence in TCAS". (Ciemier et al, 1993)

Finally, operational concerns exist with the use of TCAS. For example, in comparisons with the break-out instructions given by ATC controllers, TCAS II was found to generate avoidance maneuver commands that conflict directly with the controllers' directives in nearly 55% of the cases examined. This presents pilots with conflicting commands, and may create dangerous situations should one pilot conform to ATC commands and the other with the TCAS commands. (Toma & Massimini, 1993)

Pilot conformance to TCAS commands may also be a concern. Studies during the first years of TCAS use identified non-conformance in 24.7% of the cases reported. Of these cases, 41% were reported to be below 2500' above ground level, both immediately following take-off and during approach. Reasons given by the pilots for this non-conformance (over all phases of flight) were: (This list does not sum to 100%, as reasons were not always indicated)

• Visual Acquisition of the Other Aircraft	18.1%
• Parallel Runway Operations	16.6%
• 'Phantom Aircraft'	12.8%
• Prior ATC Communications	10.0%

Appendix E Description of the Precision Runway Monitoring System

Appendix E.1 Overview

Current independent, simultaneous approaches are limited to runways spaced 4300' apart, the minimum runway spacing in which Air Traffic Controllers are able to maintain a safe separation using the radar currently installed at major airports. The Precision Runway Monitoring (PRM) system provides faster-update, specialized radar with update rates of 1.2 or 2.4 seconds, an alerting system and specialized radar displays for an air traffic controller. PRM has been demonstrated to provide safe aircraft separation for runways as close as 3000'. This system has been tested at Memphis and Raleigh-Durham airports, and installation at several other airports is planned or underway.

Appendix E.2 PRM Alerting Logic

Incoming sensor information is monitored automatically. Two levels of automatic alerts are given to the air traffic controller. First, a warning is given when an aircraft is predicted to enter the Non-Transgression Zone (NTZ) between the parallel approaches within 10 seconds. The next level of alert is given when the aircraft enters the NTZ. As shown in Figure E.1, for a runway separation of 3400 feet, the NTZ spans a central 2000 foot-wide corridor. This configuration leaves 700 feet on either side between the NTZ and each approach path.

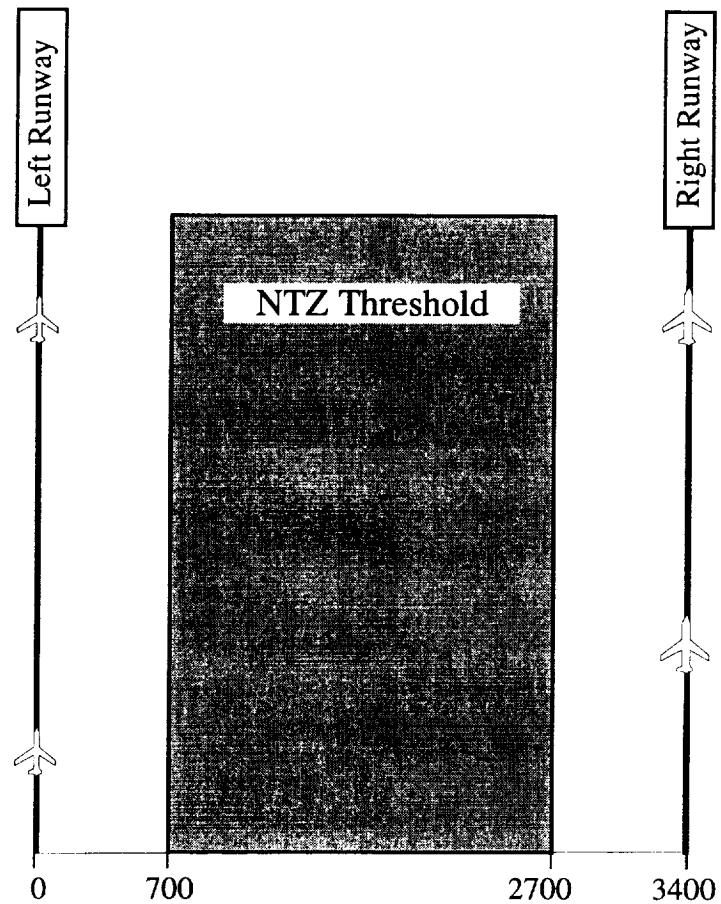


Figure E.1 Non-Transgression Zone Criteria Used by PRM

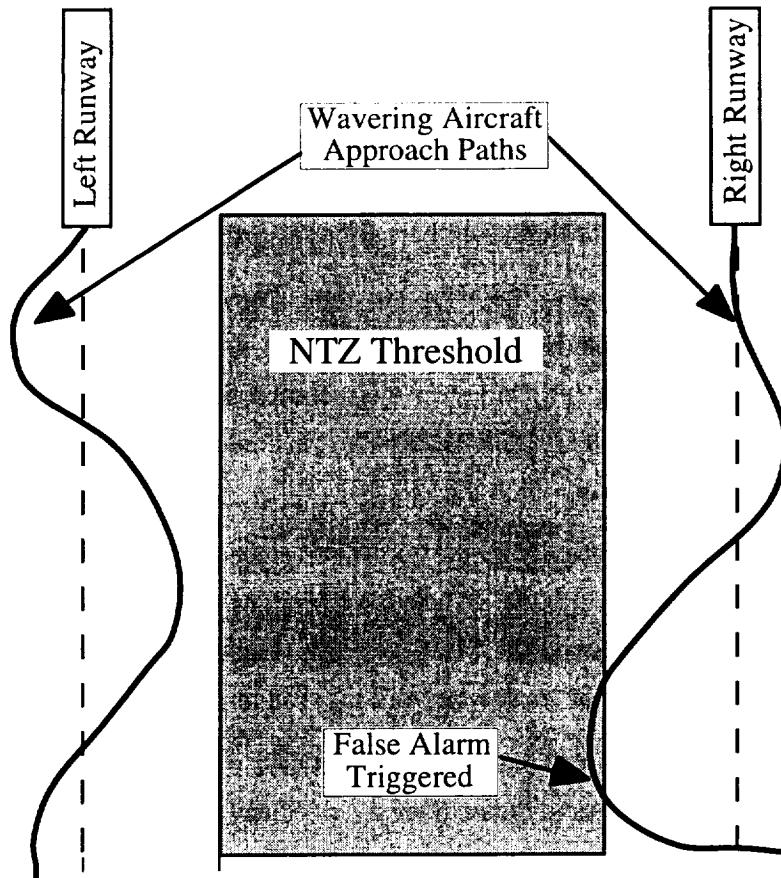


Figure E.2 Schematic of False Alarms Caused by Aircraft Deviations Into or Towards the NTZ During Nominal Approaches

Analysis of aircraft trajectories during nominal approaches has found lateral deviations about the approach course increase linearly from a standard deviation of 70 feet at the runway threshold to a standard deviation of 350 feet when the aircraft are 10 miles from the runway threshold. This lateral deviation can trigger false alarms when the aircraft deviate -- or are projected to deviate -- into the NTZ while actually maintaining their nominal approach. This effect is shown schematically in Figure E.2, and is more likely further from the runway threshold. This effect may limit the use of PRM to runways with at least 3000' separation. (Owen, 1993)

Appendix E.3 Controller Interface and Reaction Times

The PRM system is monitored by an air traffic controller. The controller's display normally covers the final approach, runway and missed approach areas, and can be expanded for a higher

resolution picture of the approaches. The runways and approach paths are displayed in white and the NTZ between the approach paths is drawn in red. Between the approach paths and the NTZ, a series of white, finer lines demarcate 100 foot increments of lateral deviation from the approach path towards the NTZ. The position of each aircraft is drawn on the screen, with alphanumeric tags providing information such as flight number, aircraft type and altitude.

When an aircraft deviates towards the NTZ, the controller is responsible for deciding upon and ordering steering commands to the deviating aircraft. These commands may steer the aircraft back towards its approach path, or they may initiate a missed approach. Additionally, if the deviating aircraft may threaten an aircraft on a parallel approach, the controller may additionally command an avoidance maneuver of that aircraft.

These commands are given verbally by the controller to the pilots using voice communication frequencies. The PRM controller has the ability to interrupt or over-ride other transmissions from other air traffic controllers. However, the PRM controller can not interrupt transmissions emanating from aircraft, but instead must wait until the frequency is open.

Figure E.3 shows the delays in each stage of this process, from the start of an aircraft's blunder to the establishment of an avoidance maneuver. Studies have examined these delays; while quick reactions are generally attained, much longer delays have occasionally been noted, with several causes. The controller may have difficulty transmitting commands to the pilots. The pilot may delay their response or use the autopilot to initiate an avoidance maneuver, which may not be aggressive enough to execute a maneuver of the required severity.

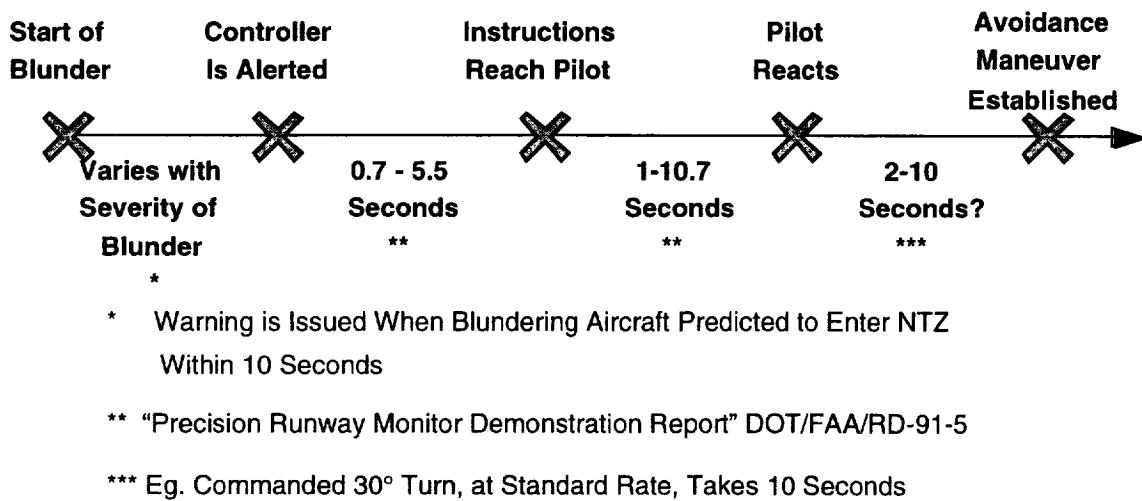


Figure E.3 Time Line of PRM Collision Avoidance Sequence

Appendix F Models of Possible Reconciliation Strategies

Chapter 6 suggested pilots may make their own evaluations of the hazard and the situation, and then attempt to reconcile their decisions with the alerting system commands. While the exact method used by the pilot for this reconciliation can not be determined, several methods can be postulated. This section will describe these methods, the difficulty the pilot may have in performing them, and the implications each would have on the resulting actions of the pilot.

A simple method of reconciliation is to ignore the alerting system's commands. This method may be used in extreme cases when the pilot has very little trust in the alerting system, or has difficulty in understanding or executing the alerting system's commands. By not requiring the pilot to monitor the alerting system or compare its results to their own reasoning, this method effectively reduces the system behavior to that achieved without an alerting system available.

A more intensive method has the pilot execute his own reasoning in parallel with the alerting system, and then generate the final commands by comparing both sets of results. In this case, the alerting system's commands become additional inputs into the pilot's reasoning about the task. This type of reconciliation may be used when the pilot does not fully trust the alerting system's logic for the task. In addition, this type of reconciliation may be used when the pilot does not understand the alerting system's functioning sufficiently to attempt to justify its results.

This model is shown schematically in Figure F.1 for the alerting sub-task; the pilot assesses the hazard level, the alerting system generates an alert, and based upon these combined values, a subsequent reconciliation process evaluates the final alert decision.

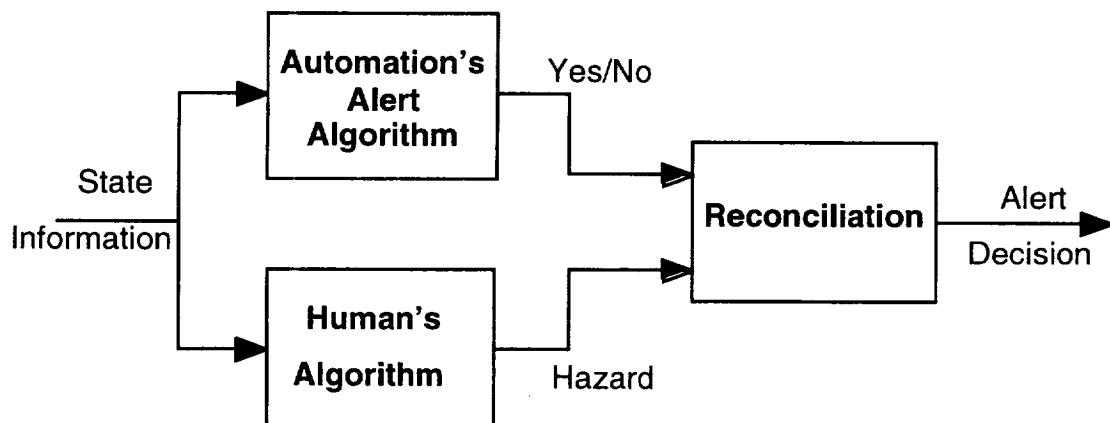


Figure F.1 Parallel Reasoning of the Pilot and Alerting System for the Alerting Sub-Task; Final Alerting Decision Based on Reconciliation of Each

Some of the subjects' comments in Experiment #3 matched this type of reconciliation. In a free responses to the question "How did the alerts (given by an automatic alerting system) change your decisions?", subjects replies indicated a willingness to alert sooner or with less certainty when automatic alerts were given in 8 of 12 responses.

Another possible method of reconciliation has the pilot execute the alerting and decision-making sub-functions after the alerting system. The alerting system acts as a trigger, and the final decision is the result of the pilot's reasoning. This method is shown for the alerting sub-task in Figure F.2. Once the pilot uses this method to assess the need for an alert, he or she may follow the commanded avoidance maneuver or also reconcile their selection of a resolution to the hazard with the automatic commands.

This method of reconciliation may be used when the alerting system is perceived to be over-conservative. For example, in the alerting sub-task, this method of reconciliation will not prevent type I errors (where the alerting system does not detect a problem), but may detect what the pilot would perceive as type II errors (where the alerting system is generating a false alarm).

This serial method of reconciliation places a lower burden on the pilot than the parallel method described before, as it only requires the pilot to perform the sub-task when the alerting system issues an alert, and, possibly, decide when the alerting system's actions are extreme enough to warrant their own solution. However, the primary benefit of the alerting system is its effect as a trigger to the pilot; once cued, the pilot's decisions alone will determine the system behavior and the pilot is still executing a larger than anticipated task load.

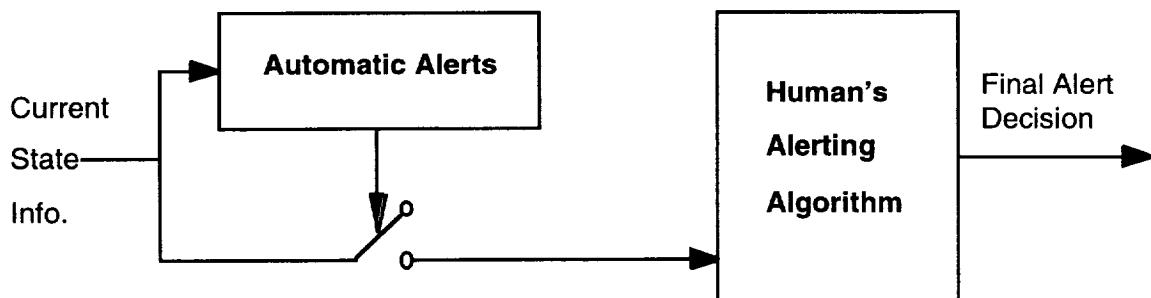


Figure F.2 Serial Method of Reconciliation

Some of the subjects' comments in Experiment #3 matched this type of reconciliation. In a free responses to the question "How did the alerts (given by an automatic alerting system) change your decisions?", subjects replies included a perception that the alerting system helped by diverting their attention from the workload inducing side-task to the collision avoidance task in 4 of the 12 responses.

A more elaborate method of reconciliation has the pilot attempt to reconstruct the alerting system's logic, as shown in Figure F.3. The final decision may follow the automatic commands or not, depending on the validity of what the pilot perceives the alerting system's rationale for its decision to be. This method of reconciliation is analogous to the 'coping strategies' found to be used by pilots when interacting with complex, difficult to understand forms of cockpit automation.

Several factors may be used to reconstruct the alerting system's logic. The pilot's independent reasoning may be used as a baseline. The inputs perceived to be used by the alerting system will be examined and any knowledge of the alerting system's logic may be applied. The pilot's perception of the alerting system's logic may be comprehensive, or it may be based upon simple heuristics.

When the pilot has a good understanding of the alerting system's functioning, the reconstruction process may provide valuable understanding of the factors underlying the decision.

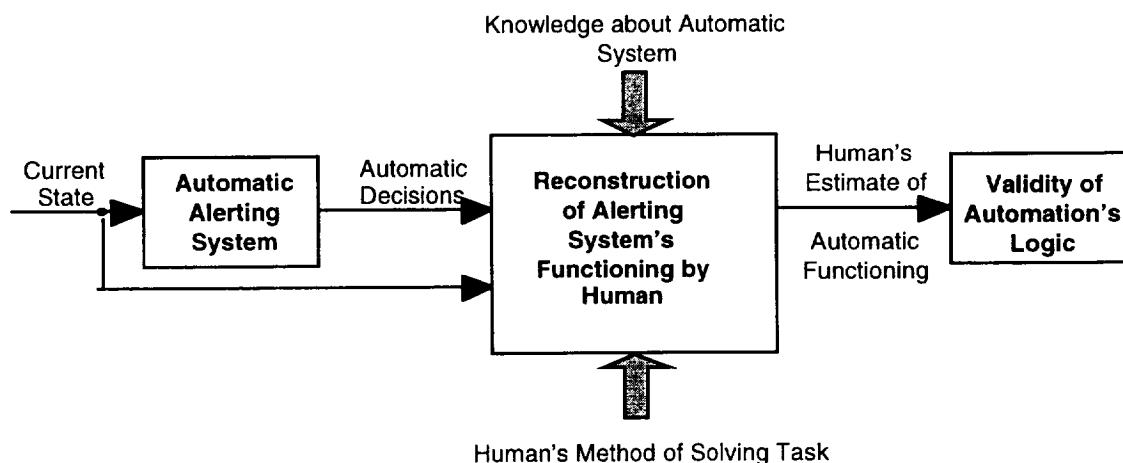


Figure F.3 Reconstruction of Alerting System's Logic

However, if the pilot does not have a good understanding of the alerting system's functioning, then the final decision may be made on an erroneous reconstruction. Therefore, this process requires knowledge of the alerting system's functioning to be accurate.

This process may also require the pilot to perform a substantial amount of computation. Several possible theories about the alerting system's reasons may need to be evaluated. Even with knowledge of the alerting system's underlying logic, this logic itself may require a large amount of computation from the pilot.

This type of reconciliation may increase or decrease the pilot's trust in the alerting system. If the alerting system's logic is understood by the pilot and he can easily reconstruct - or predict - its decisions, then the alerting system will appear to be consistent and reliable. Otherwise, the pilot may regard the alerting system as spurious and unreliable.